

# Decomposing, Comparing, and Synthesizing Access Control Expressiveness Simulations

## (Extended Version)

William C. Garrison III  
Department of Computer Science  
University of Pittsburgh  
bill@cs.pitt.edu

Adam J. Lee  
Department of Computer Science  
University of Pittsburgh  
adamlee@cs.pitt.edu

**Abstract**—Access control is fundamental to computer security, and has thus been the subject of extensive formal study. In particular, *relative expressiveness analysis* techniques have used formal mappings called *simulations* to explore whether one access control system is capable of emulating another, thereby comparing the expressive power of these systems. Unfortunately, the notions of expressiveness simulation that have been explored vary widely, which makes it difficult to compare results in the literature, and even leads to apparent contradictions between results. Furthermore, some notions of expressiveness simulation make use of non-determinism, and thus cannot be used to define mappings between access control systems that are useful in practical scenarios. In this work, we define the minimum set of properties for an *implementable* access control simulation; i.e., a deterministic “recipe” for using one system in place of another. We then define a wide range of properties spread across several dimensions that can be enforced on top of this minimum definition. These properties define a taxonomy that can be used to separate and compare existing notions of access control simulation, many of which were previously incomparable. We position existing notions of simulation within our properties lattice by formally proving each simulation’s equivalence to a corresponding set of properties. Lastly, we take steps towards bridging the gap between theory and practice by exploring the systems implications of points within our properties lattice. This shows that relative expressive analysis is more than just a theoretical tool, and can also guide the choice of the most suitable access control system for a specific application or scenario.

### I. INTRODUCTION

Access control is foundational to computer security and, as such, has been the topic of extensive formal study. Much of this work has focused on comparing different techniques for representing and enforcing access control, deemed access control *models*, *systems*, or *schemes*. By far the most common type of comparative study in access control techniques is the expressiveness simulation (e.g., [1]–[14]). A simulation is a formal mapping from, say, system  $S$  to system  $T$  that proves  $T$  is at least as expressive as  $S$ : that is,  $T$  possesses the raw capability to be used in operating environments in place of  $S$ .

However, the formal definitions of the various simulations used in the literature vary widely. Different simulations have been used to prove various types of results, ranging from very specific properties about whole ranges of models (e.g., monotonic access control models with multi-parent creation

cannot be simulated by monotonic models with only single-parent creation [6]) to the ability to replace certain specific models with others in practice (e.g., role-based access control can be configured to enforce mandatory and discretionary policies [9]). However, this disparity in the goals of these works has led to many different definitions of access control simulation, often tailored to the particular result sought. It has been shown that these different simulations prove wildly different notions of expressiveness, often not preserving any particular security properties [13].

Furthermore, not all of these notions of simulation are practically useful. For instance, some make use of non-determinism, manipulating the policy differently depending on what future queries will be asked. While this may allow a theorist to show that system  $T$  is capable of doing all the things  $S$  is, if a practitioner wants to use system  $T$  in place of system  $S$ , she needs a deterministic procedure for doing so.

In this work, we build a taxonomy for expressiveness simulations based on the simulation properties that they satisfy. We determine the minimum requirements for a mapping to be *implementable*, or applicable toward using one system in place of another in practice. We use these requirements to construct a general definition of implementable simulation, and provide a taxonomy of additional restrictions on this definition for simulations that enforce more stringent properties. We then position existing simulations from the literature within this lattice, providing the first such comparison in the literature.

To this end, we make the following contributions.

**Definition of implementable access control simulation** We propose a general definition of an *implementable* access control mapping that is broad enough to encompass much of the wide range of existing access control simulations, yet precise enough to guarantee implementability. Intuitively, an *implementable* simulation of  $S$  in  $T$  shows that  $T$  can accomplish everything  $S$  can, and deterministically shows *how* (Section III).

**Lattice of simulation properties** We decompose and expand upon the properties enforced by various access control simulations from the literature, forming a lattice relating the range of access control simulations to one another. This lattice allows us to formally compare the guarantees offered by

existing notions of access control simulation (many of which were not formerly known to be comparable) and points to unexplored combinations of properties that can yield different expressiveness results (Section IV).

**Positioning of existing simulations** We construct formal proofs positioning existing notions of access control simulation within our lattice of simulation properties, including a comparative discussion of simulations that previously seemed incomparable. We thus systematize the formal relationships between previously-published simulations, allowing reconciliation of previously disparate expressiveness knowledge (Section V).

**Selecting simulation properties** We observe that many of the dimensions upon which our simulation property lattice is built have implications for the use of simulations for satisfying real-world requirements using existing access control systems (e.g., required storage, whether data structures must be locked for concurrent usage). Thus, in addition to positioning existing notions of simulation within our lattice of properties, we assist in creating new notions of simulation by selecting the properties that should be enforced in an expressiveness analysis based upon the scenario in which an eventual access control deployment will occur. To this end, we discuss in detail various interactions between simulation properties, the results of enforcing different properties, and how a specific deployment scenario dictates which properties are relevant (Section VI).

We begin by providing background on the goals of and techniques used in relative expressiveness analysis.

## II. RELATIVE EXPRESSIVENESS ANALYSIS

In this section, we describe how relative expressiveness analysis is conducted, survey the history of the technique, and point out the wide variety in existing access control expressiveness simulations.

### A. Motivating Examples

An access control system’s *expressiveness* (or expressive power) is a measure of the range of policies that it can represent and the transformations it can make to those policies. Statements of *relative* expressiveness state that one system is capable of replacing another (that is, it can represent all the same policies and transform them in equivalent ways). Assume, for instance, that an organization is considering transitioning from one access control solution to another, in order to accommodate evolving requirements. The organization may have specific desired features for this new access control system, but it certainly must be able to represent all of the policies that the existing system can, or it would not be a suitable replacement. Thus, this organization is searching for a new system that is *at least as expressive as* its old system.

Another use of relative expressiveness is in suitability analysis. Prior work has noted that *practically* evaluating an access control system must take into account the application in which the system is to be used, as well as additional cost metrics (e.g., computation, ease of use). This analysis problem has been identified as a system’s *suitability* to a particular application [15], [16]. Suitability analysis formalizes

an application’s access control requirements (a *workload*), and uses expressiveness to prove that an access control system can satisfy those requirements. Assume, in this case, that the aforementioned organization is choosing an initial access control system for a new collection of data. Comparing the candidates’ relative expressiveness is not particularly enlightening, since the most expressive system may not be the most suitable; the organization should instead formalize their access control workload and use relative expressiveness analysis to identify which of the candidates are *expressive enough* to satisfy this workload. Thus, while work in suitability analysis has shown that expressive power alone is insufficient for evaluating an access control system, expressiveness is a fundamentally important component of a more general suitability analysis workflow: one cannot determine which access control system is *best* for a particular use case without first determining which are *capable of satisfying* that use case.

### B. Prior Work

Relative expressiveness analysis generally starts by formalizing a pair of access control systems as state machines. These state machines include, at a minimum: a set of states, each of which encapsulates a snapshot of the access control system’s data structures; a procedure describing how to interpret the states’ data structures to determine which authorization requests are granted; and a set of commands, used to manipulate the data structures and thus transition between states. Some formalisms for access control systems also include additional queries beyond access requests [13], [14]. A *simulation*, then, is a structure that proves  $\mathcal{T}$  is at least as expressive as  $\mathcal{S}$ —or, that  $\mathcal{T}$  can be used in place of  $\mathcal{S}$ . The term *simulation* is rather vague, here, and for good reason: various notions of simulation in the literature have meant very different things (e.g., What type of behavior must be simulated? How closely must  $\mathcal{T}$  represent the information in  $\mathcal{S}$ ?), and as a result have implied very different types of expressiveness results.

The works of Sandhu, Ganta, Munawer, and Osborn [2]–[4], [7]–[9] include some of the earliest access control simulations. In these works, a simulation of  $\mathcal{S}$  in  $\mathcal{T}$  must show that a permission can be granted in  $\mathcal{S}$  if and only if it can also be granted in  $\mathcal{T}$ . No other formal properties are enforced, though in some cases additional properties become part of the *de facto* definition of simulation. For instance, while there is no requirement for  $\mathcal{T}$  to have a state equivalent to each  $\mathcal{S}$  state (merely for  $\mathcal{T}$  to be able to grant each access that  $\mathcal{S}$  does, in some state), the example simulations all include methods for mapping each  $\mathcal{S}$  state to a  $\mathcal{T}$  state (as this is the simplest way to show the required property). In addition, although the definition does not prohibit the use of an unbounded number of  $\mathcal{T}$  commands to simulate a single  $\mathcal{S}$  command, Sandhu and Munawer [7] only use simulations in which an  $\mathcal{S}$  command is simulated using a constant number of  $\mathcal{T}$  commands.

Ganta’s PhD dissertation [5] attempts to formalize a more rigorous notion of expressiveness simulation. In his simulation, the state correspondence is explicit, requiring that each state in  $\mathcal{S}$  have a corresponding state in  $\mathcal{T}$  that grants all the same

accesses (at least, all those that exist in  $\mathcal{S}$ —those that exist in  $\mathcal{T}$  but not in  $\mathcal{S}$  are unconstrained). In addition, to ensure that  $\mathcal{T}$  cannot grant accesses that  $\mathcal{S}$  cannot, any state that can be entered in  $\mathcal{T}$  must also have a corresponding reachable state in  $\mathcal{S}$ . Finally, to ensure accesses in  $\mathcal{T}$  cannot be combined in ways that cannot occur in  $\mathcal{S}$ , the following restriction is made: when simulating a  $\mathcal{T}$  command in  $\mathcal{S}$ , multiple commands may be used, but each state along the way must allow *either* a subset of the accesses of the start state or a subset of the accesses of the end state. Thus, no two accesses can be allowed in the same state in  $\mathcal{T}$  that are not allowed in a single state in  $\mathcal{S}$ .

Ammann, Lipton, and Sandhu [1], [6] took a different (and much more strict) approach to more rigorously defining a simulation. First, they describe a strict state correspondence that requires  $\mathcal{T}$  to represent its states with the same sets and relations as  $\mathcal{S}$ , and for these sets to have identical contents in corresponding  $\mathcal{T}$  and  $\mathcal{S}$  states. In other words,  $\mathcal{T}$  cannot include additional elements in any sets that  $\mathcal{S}$  uses (although additional, distinct sets may be stored). For example, one could simulate the state  $\{U = \{a, b\}, V = \{c\}\}$  with state  $\{U = \{a, b\}, V = \{c\}, W = \{\langle a, d \rangle, \langle b, d \rangle\}\}$ , but not with  $\{U = \{a, b\}, V = \{c, d\}\}$ . Given this notion of state correspondence, a simulation then shows that  $\mathcal{T}$  can reach a state corresponding to each reachable  $\mathcal{S}$  state, and cannot reach any state that does not have a reachable corresponding state in  $\mathcal{S}$ . This strict notion of simulation is used to show that monotonic, multi-parent systems are more expressive than monotonic, single-parent systems (e.g., there are monotonic multi-parent systems that cannot be simulated by any monotonic single-parent system).

Chander, Dean, and Mitchell [10] restrict the definition of simulation in a different way. Rather than force a more strict state correspondence (the static portion of the simulation), they more tightly restrict the way the simulation handles the system as it executes (i.e., the command mapping). In these simulations, the state correspondence is comparatively lax: to simulate an  $\mathcal{S}$  state, a  $\mathcal{T}$  state must allow and deny all the same authorization requests as its corresponding  $\mathcal{S}$  state. Additional requests can exist in  $\mathcal{T}$  and are unconstrained, but all requests corresponding to those in  $\mathcal{S}$  must have the same value in corresponding states. However, the process for simulating an  $\mathcal{S}$  command using  $\mathcal{T}$  commands must be independent of the state: it cannot execute a  $\mathcal{T}$  command for each user, or otherwise inspect the state when determining what commands should be executed. In addition, in the strong form of simulation, each  $\mathcal{S}$  command must be simulated with a single  $\mathcal{T}$  command. They then compare the expressiveness of access control lists, trust management, and two forms of capability systems (all systems studied in forms with and without revocation and delegation).

Tripunitara and Li [12], [13] noted that the existing notions of simulation did not correspond directly to any particular safety analysis questions, and thus a simulation of any of these types does not make any particular safety guarantees. They formalize *compositional security analysis* (intuitively, determining whether a certain set of access control queries will always, never, or sometimes become true in any reachable state), which is a generalization of simple safety analysis [17].

They then present a notion of simulation tailor-made to preserve these types of analysis questions.

Their simulation, called the state-matching reduction, considers a broader range of queries than only authorization requests, placing the strictness of its state correspondence somewhere between the work of Ammann, Lipton, and Sandhu and that of Chander, Dean, and Mitchell. The state-matching reduction maps each query  $q^{\mathcal{S}}$  in  $\mathcal{S}$  to a single query  $q^{\mathcal{T}}$  in  $\mathcal{T}$ , and the simulation must determine the value of  $q^{\mathcal{S}}$  in any state in  $\mathcal{T}$  by checking the value of  $q^{\mathcal{T}}$ . Finally, reachability constraints ensure that  $\mathcal{T}$  can reach a state corresponding to each reachable  $\mathcal{S}$  state, and cannot reach any state that does not have a reachable corresponding state in  $\mathcal{S}$ . Tripunitara and Li prove that this notion of simulation preserves compositional security analysis instances: that is, if there exists a state-matching reduction from  $\mathcal{S}$  to  $\mathcal{T}$ , then any compositional security analysis instance has the same truth value in both systems. Tripunitara and Li’s reductions have since been used to analyze role-based access control [18] and prove that newly-proposed systems are more expressive than certain existing systems [19].

Work by Hinrichs et al. [14] recognizes the value of the state-matching reduction but claims that, in practice, not all scenarios require the preservation of all possible compositional security analysis instances (nor are these the only types of safety properties that are ever relevant). They present *parameterized expressiveness*, which defines a baseline set of simulation properties, and provides several additional properties that can be enforced atop the baseline to provide additional guarantees. The base simulation uses the same query-based state correspondence as Tripunitara and Li, but relaxes the query mapping to allow it to consult multiple  $\mathcal{T}$  queries to determine the value of an  $\mathcal{S}$  query during simulation. Further properties enforced above this baseline include using the identity query mapping for authorization requests (to ensure that  $\mathcal{T}$ ’s authorization questions are the queries being used to simulate  $\mathcal{S}$ ’s authorization requests), forbidding string manipulations (to prohibit the state mapping from using arbitrary encodings to store information in the contents of strings such as user names), and restricting the command mapping from mapping non-administrative commands in  $\mathcal{S}$  to administrative commands in  $\mathcal{T}$ . This framework has since been used to evaluate the suitability of certain general-purpose access control systems for various unique, application-specific requirements [15], [16].

### C. Usage and Implications

Unfortunately, there are several indications that research on expressiveness analysis is being held back by the inability to reconcile the vastly different notions of expressiveness simulations and the disconnect between the properties preserved by a simulation and those that are important to a practical deployment. Several works have demonstrated scenarios in which static notions of expressiveness indicate two systems are equally capable of satisfying a set of operational requirements, but in practice they are better-suited to very different deployment scenarios [15], [20]. Bourdier et al. point out the existence of several competing techniques for expressiveness analysis,

none of which consider the deployment. They approach one facet of this problem by proposing a formalism for access control systems that can more easily be transformed into implementations using rewrite-based tools [21]. Several others simply express a desire to use expressiveness analysis, but never do so, presumably due to the complexities of selecting and using the right notion of simulation [22], [23].

A group at the National Institute of Standards and Technology has developed Policy Machine, an attempt at a universal access control system (one that can represent any policy via only configuration changes) [24]. However, in evaluating Policy Machine’s success, they avoid formally proving its expressiveness and instead show informal mappings that demonstrate how one might use Policy Machine to represent several existing access control systems’ policies [25]. Soon after, the group published a report bemoaning the lack of quality metrics for evaluating access control systems, noting that, in access control, “one size does not fit all,” and thus said metrics must consider the deployment scenario [26].

This overview illustrates that while each notion of expressiveness simulation has been used to prove various results, the body of knowledge is troublesome to interpret and utilize due to the wide variation in the properties required by each simulation. In this work, we fill this void in the literature by (1) proposing a minimal definition of simulation that satisfies properties guaranteeing that its results are practically useful; (2) presenting a set of additional properties that more strict simulations can enforce; and (3) categorizing the above notions of simulation based on the properties that they enforce. We make the additional contribution of (4) discussing relationships between attributes of a deployment scenario and the practical effects of enforcing simulation properties, thus assisting analysts in selecting the most relevant properties (and therefore conducting the most relevant form of expressiveness analysis) for the environment in which an access control system will be deployed.

### III. IMPLEMENTABLE EXPRESSIVENESS SIMULATIONS

In this section, we give requirements for a simulation to be *implementable* and define our general formulation of relative expressiveness analysis through the lens of implementability.

#### A. Implementability Requirements

In this work, we aim to consider expressiveness simulations that are *implementable*: i.e., practically useful for making decisions about which system is most suitable for a particular deployment. Implementability enforces the following intuition: if a system  $\mathcal{T}$  is at least as expressive as  $\mathcal{S}$ , then one should be able to determine a general way to use  $\mathcal{T}$  in place of  $\mathcal{S}$ . Thus, we define a minimal set of properties for an expressiveness mapping to be considered implementable.

**State mapping** In order to use  $\mathcal{T}$  in place of  $\mathcal{S}$ , it must be possible to (uniquely) determine which  $\mathcal{T}$  state to use in place of a particular  $\mathcal{S}$  state. Thus, the state mapping must be a

function from the simulated system states to the simulating system states.<sup>1</sup>

**Command mapping** To use  $\mathcal{T}$  in place of  $\mathcal{S}$ , it must be possible to execute commands in  $\mathcal{T}$  that are equivalent to the commands in  $\mathcal{S}$ . It is not necessarily the case that each  $\mathcal{S}$  command can be simulated using a single  $\mathcal{T}$  command, so we require a function from  $\mathcal{S}$  commands to *sequences of  $\mathcal{T}$  commands*.<sup>2</sup> Finally, it may be necessary to map an  $\mathcal{S}$  command differently depending upon the state in which it is intended to be executed. Since using  $\mathcal{T}$  in place of  $\mathcal{S}$  means we only have a  $\mathcal{T}$  state to inspect during execution, this function should map an  $\mathcal{S}$  command and a  $\mathcal{T}$  state to a sequence of  $\mathcal{T}$  commands.

**Query decider** For some simulations of  $\mathcal{S}$  in  $\mathcal{T}$ , we may only care that  $\mathcal{T}$  allows the same set of accesses that  $\mathcal{S}$  would. However some types of simulations may allow the overriding of  $\mathcal{T}$ ’s default method of deciding granted permissions (e.g., adding the additional requirement that the requesting user is a member of the `REAL_USERS` group, to distinguish from other data stored in the user-set). While some types of simulations do not allow this, to remain general we simply require a function that maps each  $\mathcal{S}$  query and  $\mathcal{T}$  state to either true or false. In some formalisms, this only includes the queries requesting access, while in other cases other types of queries are allowed (e.g., “Is user  $u$  a member of role  $r$ ?”).

We use these requirements to motivate our definition of the general case of implementable relative expressiveness.

#### B. Expressiveness Mappings

To define relative expressiveness mappings, we must first define the state machines that represent access control systems. Since we aim to compare existing expressiveness simulations, we use a formalism for these structures that remains similar to existing work, e.g., [10], [12], [13].

An access control system is formalized as a state machine belonging to a particular access control *model*. An access control model formalizes the way in which the access control system will store and interpret information to make access control decisions. Its data structures are formalized as a set of access control *states*, and its methods for determining whether to allow or deny inquiries as a set of *authorization requests*. The value of all requests in a state (whether they are allowed or denied) defines the access control policy, or *theory*, to be enforced in that state.

**Definition 1** An access control model is defined as  $\mathcal{M} = \langle \Gamma, \mathcal{R} \rangle$ , where  $\Gamma$  is the set of states and  $\mathcal{R}$  is the set of authorization requests, where each request  $r \in \mathcal{R}$  is a function  $\Gamma \rightarrow \{\text{TRUE}, \text{FALSE}\}$ . The entailment ( $\vdash$ ) of a request is defined as  $\gamma \vdash r \triangleq r(\gamma) = \text{TRUE}$ .  $\diamond$

<sup>1</sup>It is possible that multiple states in  $\mathcal{S}$  can be represented using the same state in  $\mathcal{T}$ . Thus, we do not require the state mapping to be an injection. Furthermore, there may be states in  $\mathcal{T}$  that are not used to simulate  $\mathcal{S}$ , and thus the state mapping need not be a surjection.

<sup>2</sup>Not *sets* of  $\mathcal{T}$  commands, as commands may appear multiple times; and not *bags* of  $\mathcal{T}$  commands, as order matters.

For example, consider a simple role-based access control model whose states are defined over sets  $U$  of users,  $P$  of permissions, and  $R$  of roles, as well as the user assignment  $UR \subseteq U \times R$  and permission assignment  $PA \subseteq R \times P$ . The requests in this model are of the form “Is  $u$  authorized for  $p$ ?,” which is TRUE if  $\exists r : \langle u, r \rangle \in UR \wedge \langle r, p \rangle \in PA$ .

When we refer to the *size* of a state, we are referring to the size of its decomposition into primitive objects (e.g., users and roles) and tuples (e.g., entries in a user assignment relation).

**Definition 2** Given an access control model  $\mathcal{M} = \langle \Gamma, \mathcal{R} \rangle$  and a state  $\gamma \in \Gamma$ , the set decomposition of  $\gamma$  is denoted  $[\gamma]$ , and refers to the “set of sets” forming  $\gamma$ , in which  $\gamma$  is represented as being comprised of primitive sets and relations.  $\diamond$

Thus, the size of an access control state  $\gamma$  is defined as  $|\gamma| = \sum_{S \in [\gamma]} |S|$ . For example, if  $[\gamma] = \{U = \{u_1\}, R = \{r_1, r_2\}, UR = \{\langle u_1, r_1 \rangle, \langle u_1, r_2 \rangle\}\}$ , then  $|\gamma| = |U| + |R| + |UR| = 5$ .

An access control system expands on a model by providing methods of transforming the current state and additional methods of querying the states. These additional queries allow the user to ask additional boolean queries of the system, but a value of TRUE does not indicate an authorization was granted.

**Definition 3** Given access control model  $\mathcal{M} = \langle \Gamma, \mathcal{R} \rangle$ , an access control system within  $\mathcal{M}$  is a state transition system,  $\mathcal{S} = \langle \Gamma, \Psi, Q \rangle$ , where  $\Psi$  is the set of commands, where each command  $\psi \in \Psi$  is a function  $\Gamma \rightarrow \Gamma$ , and  $Q \supseteq \mathcal{R}$  is the set of queries, where each query  $q \in Q$  is a function  $\Gamma \rightarrow \{\text{TRUE}, \text{FALSE}\}$ .  $\diamond$

We use the notation  $\text{next}(\gamma, \psi)$  to denote the state resulting from executing  $\psi$  in  $\gamma$  (that is,  $\psi(\gamma)$ ), and  $\text{terminal}(\gamma, \psi_1 \circ \psi_n)$  to denote the final state produced by repeatedly applying  $\text{next}$  to the commands  $\psi_1, \dots, \psi_n$  starting from state  $\gamma$ :  $\text{next}(\dots \text{next}(\gamma, \psi_1), \dots, \psi_n)$ .

A system based on the example role-based model must define commands to transform the state: e.g., to assign roles to users, and assign permissions to roles. Additional queries beyond the model’s requests may include those of the form “Is user  $u$  a member of role  $r$ ?”

Next, we define an access control mapping, which maps one system to another but does not enforce any simulation properties. We define a mapping as motivated in Section III-A so that it can represent any implementable expressiveness simulation.

**Definition 4** Given two access control systems,  $\mathcal{S} = \langle \Gamma^S, \Psi^S, Q^S \rangle$  and  $\mathcal{T} = \langle \Gamma^T, \Psi^T, Q^T \rangle$ , a mapping from  $\mathcal{S}$  to  $\mathcal{T}$  is a triple of functions  $\sigma = \langle \sigma_\Gamma, \sigma_\Psi, \sigma_Q \rangle$ , where:

- $\sigma_\Gamma : \Gamma^S \rightarrow \Gamma^T$  is the state mapping
- $\sigma_\Psi : \Psi^S \times \Gamma^T \rightarrow (\Psi^T)^*$  is the command mapping
- $\sigma_Q = Q^S \times \Gamma^T \rightarrow \{\text{TRUE}, \text{FALSE}\}$  is the query decider

This definition is demonstrated in Fig. 1. Each function takes its most general form that satisfies the requirements from Section III-A. Thus, the definition remains general (it does not enforce any specific security requirements yet), while




Fig. 1: The general form of an implementable expressiveness mapping.

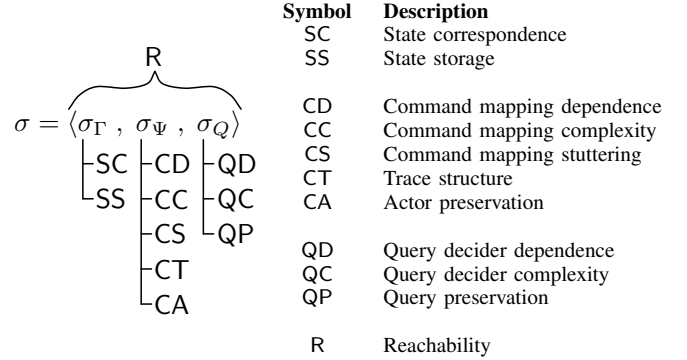


Fig. 2: An overview of the dimensions of expressiveness simulation properties

ensuring that any such mappings can generate implementable procedures for using the simulating system in place of the simulated system.

To demonstrate Definition 4, consider mapping a simple access control list system to the role-based system described throughout this section. The state mapping can map each ACL state to a role-based state in which each user  $u$  has a unique role  $r_u$ , and each user’s role is assigned the permissions from the ACL state. The command mapping can map, e.g., “grant  $u$  access to  $o$ ” to “assign  $o$  to role  $r_u$ .” The query mapping would then map “Can  $u$  access  $o$ ?” to “Is  $u$  authorized for  $o$ ?”

#### IV. EXPRESSIVENESS SIMULATION PROPERTIES

In this section, we describe the lattice of properties that we use to taxonomize access control expressiveness simulations.

##### A. Overview of dimensions of properties

In order for a mapping to be considered a simulation, it must enforce additional properties over Definition 4. We restrict this definition in a variety of ways. Although no set of restrictions can be shown to be the full, correct set for all conceivable simulations, there are naturally three categories of restrictions to consider for simulations, given their structure (a set of three functions): i.e., refinements to each of the state correspondence, command mapping, and query decider. We also consider restrictions to the reachability constraints required (a cross-cutting dimension describing how these functions must relate to one another). A summary of these dimensions is depicted in Fig. 2.

Our state correspondence  $\sigma_\Gamma$  can be based on any of a handful of structural definitions, defined by SC (i.e., what elements do we inspect to determine whether two states correspond?). Further, SS can limit the amount of storage the state correspondence uses (e.g.,  $\mathcal{T}$  must simulate  $\mathcal{S}$  using only a linear amount of additional storage).

The command mapping  $\sigma_\Psi$  can be restricted by CD in what state elements it can use to map commands (e.g., whether it can inspect arbitrary state elements or only those that are exposed via queries). CC considers limiting the time-complexity of the command mapping routine. Since the command mapping returns a *sequence* of commands, CS can limit the number of commands it can return (e.g., only one, or constant in the size of the state). We identify CT, a dimension of concurrency-related trace structure restrictions, as well as CA, requiring the simulation to map  $\mathcal{S}$  commands executed by certain types of users only to  $\mathcal{T}$  commands executed by that category of users.

The query decider  $\sigma_Q$  can also be restricted in a number of ways. Like the command mapping, we may limit what elements of the state the decider can inspect when deciding how to answer queries within a specific state (QD), or the time-complexity of the routine (QC). In some cases a simulation of  $\mathcal{S}$  in  $\mathcal{T}$  may be required to map certain  $\mathcal{S}$  queries to specific related queries in  $\mathcal{T}$ , most notably authorization requests (e.g., to answer whether user  $u$  should have permission  $p$  in  $\mathcal{S}$ ,  $\mathcal{T}$  should simply check whether user  $u$  has permission  $p$  in its current  $\mathcal{T}$  state); this type of restriction is handled in QP.

Finally, our reachability restrictions  $R$  define how these three functions relate, by allowing us to parameterize whether we require one-way reachability ( $\mathcal{T}$  must be able to transition to states corresponding to all reachable  $\mathcal{S}$  states) or bidirectional reachability ( $\mathcal{T}$  also cannot transition to states that do not correspond to reachable  $\mathcal{S}$  states).

The bare minimum set of these simulation properties that must be enforced for a mapping to be considered a simulation is a notion of state correspondence and a reachability relation. We present the definition of implementable expressiveness simulation, which refines the mapping by enforcing these properties.

**Definition 5** Given two access control systems,  $\mathcal{S} = \langle \Gamma^S, \Psi^S, Q^S \rangle$  and  $\mathcal{T} = \langle \Gamma^T, \Psi^T, Q^T \rangle$  and a mapping  $\sigma = \langle \sigma_\Gamma, \sigma_\Psi, \sigma_Q \rangle$  from  $\mathcal{S}$  to  $\mathcal{T}$ , an implementable expressiveness simulation of  $\mathcal{S}$  in  $\mathcal{T}$  based on  $\sigma$  is defined as  $\sigma' = \langle \sigma_\Gamma, \sigma_\Psi, \sigma_Q, \sim, R \rangle$ , where:

- $\sim \subseteq \Gamma^S \times \Gamma^T$  is the state correspondence, and  $\forall \gamma \in \Gamma^S, \gamma \sim \sigma_\Gamma(\gamma)$
- $R$  is a reachability restriction

We define all properties over the expressiveness simulation  $\sigma = \langle \sigma_\Gamma, \sigma_\Psi, \sigma_Q, \sim, R \rangle$ . Unless otherwise noted, properties within a dimension are totally ordered from most to least strict.

### B. State correspondence properties

As discussed in Section III-B, the state correspondence of an implementable simulation of  $\mathcal{S}$  in  $\mathcal{T}$  is a function,  $\sigma_\Gamma$  :

$\Gamma^S \rightarrow \Gamma^T$  mapping each state in  $\mathcal{S}$  to a state in  $\mathcal{T}$ . There are several ways in which we can restrict this mapping.

### Dimension SC: State correspondence structure

This dimension of properties restricts the way in which corresponding states are structurally similar. All properties within this dimension were inspired by state correspondence relations from prior expressiveness simulations; other application-specific state correspondence relations are conceivable.

#### SCs: Structure-correspondent

$$\gamma^S \approx \gamma^T \triangleq \forall S_i \in [\gamma^S]. (S_i \in [\gamma^T])$$

#### SCq: Query-correspondent

$$\gamma^S \approx \gamma^T \triangleq \forall q \in Q^S. (\gamma^S \vdash q \iff \sigma_Q(q, \gamma^T) = \text{TRUE})$$

#### SCa: Authorization-correspondent

$$\gamma^S \approx \gamma^T \triangleq \forall r \in \mathcal{R}^S. (\gamma^S \vdash r \iff \sigma_Q(r, \gamma^T) = \text{TRUE})$$

*Authorization-correspondent* simulations enforce that every  $\gamma^S$  maps to a  $\gamma^T$  that agrees on all authorization requests: any permission granted/denied in  $\gamma^S$  must also be granted/denied in  $\gamma^T$ . Requests that exist in  $\mathcal{T}$  but not in  $\mathcal{S}$  are not restricted. This type of correspondence is used in [2], [3], [7]–[10]. *Query-correspondence* requires that  $\gamma^S$  and  $\gamma^T$  agree on *all* queries, not just authorization requests. This type of correspondence is used in the expressiveness simulations of [13], [14].

Finally, *structure-correspondent* simulations require all corresponding state elements to be identical. If  $\gamma^S$  structure-corresponds to  $\gamma^T$ , then every set in  $\gamma^S$  exists in  $\gamma^T$ , and contains all the same elements ( $\gamma^T$  may contain additional sets or relations). Thus, if  $\gamma^S$  contains sets of users and permissions, and a relation between them (a subset of users  $\times$  permissions) specifying accesses,  $\gamma^T$  must contain identical sets of users and permissions, and an identical set of  $\langle \text{user}, \text{permission} \rangle$  pairs. This notion of state correspondence is used in [6].

The type of state correspondence used is a central characteristic of a type of simulation. Enforcing a state correspondence that is too weak can allow the simulating system to diverge from the simulated system in unexpected ways, while a state correspondence that is too strong will cause the simulating system to track the simulated system more closely than necessary (e.g., by constraining the values of queries that the deployment never needs to ask). Thus, choosing a particular state correspondence is choosing how closely the simulating system must stay to the simulated system.

### Dimension SS: State storage

An orthogonal class of restrictions that can be placed on the state correspondence relation involve its allowed storage. Here, we restrict the size of  $\gamma^T = \sigma_\Gamma(\gamma^S)$  with respect to  $\gamma^S$ .

#### SSI: Linear storage

$$\exists c \in \mathbb{R}^+, s \in \mathbb{Z}^+ : \forall \gamma \in \Gamma^S : |\gamma| \geq s \Rightarrow |\sigma_\Gamma(\gamma)| \leq c|\gamma|$$

#### SSp: Polynomial storage

$$\exists k \in \mathbb{R}^+, s \in \mathbb{Z}^+ : \forall \gamma \in \Gamma^S : |\gamma| \geq s \Rightarrow |\sigma_\Gamma(\gamma)| \leq |\gamma|^k$$

#### SS $\infty$ : Unbounded storage No restriction.

A *linear storage* simulation says that  $\gamma^{\mathcal{T}}$  can grow at most linearly with  $\gamma^{\mathcal{S}}$ , while in a *polynomial storage* simulation, the size of  $\gamma^{\mathcal{T}}$  is bounded by a polynomial in the size of  $\gamma^{\mathcal{S}}$ . The most obvious result of enforcing properties within SS is limited trusted storage, but it can also limit iteration over the resulting state (e.g., if an action must be taken for each document in the simulating system, SSI ensures that this sequence of actions is linear in the size of the simulated state).

### C. Command mapping properties

Recall that the command mapping for an implementable simulation (Definition 5) is a function  $\sigma_{\Psi} : \Psi^{\mathcal{S}} \times \Gamma^{\mathcal{T}} \rightarrow (\Psi^{\mathcal{T}})^*$  that returns the sequence of  $\mathcal{T}$  commands needed to simulate an  $\mathcal{S}$  command starting from a particular  $\mathcal{T}$  state. Thus, it allows us to simulate  $\mathcal{S}$  commands in an active simulation using  $\mathcal{T}$ . We now discuss the ways in which we can restrict this mapping.

#### Dimension CD: Command mapping dependence

While Definition 5 maps each  $\mathcal{S}$  command and  $\mathcal{T}$  state to a sequence of  $\mathcal{T}$  commands, some previous works use more strict command mappings, mapping each  $\mathcal{S}$  command to a sequence of  $\mathcal{T}$  commands without considering the state [10]. In between these options, we may map each  $\mathcal{S}$  command and  $\mathcal{T}$  theory, calculating the sequence of  $\mathcal{T}$  commands by observing only the queriable portions of the  $\mathcal{T}$  state. *Command mapping dependence* thus restricts the information that the command mapping can consider about a  $\mathcal{T}$  state when calculating the trace of  $\mathcal{T}$  commands to execute.

##### CDi: Independent command mapping

$$\exists \sigma' : \Psi^{\mathcal{S}} \rightarrow (\Psi^{\mathcal{T}})^*. (\sigma_{\Psi}(\psi, \gamma) \equiv \sigma'(\psi))$$

##### CDt: Theory-dependent command mapping

$$\exists \sigma' : \Psi^{\mathcal{S}} \times Th(\mathcal{T}) \rightarrow (\Psi^{\mathcal{T}})^*. (\sigma_{\Psi}(\psi, \gamma) \equiv \sigma'(\psi, Th(\gamma)))$$

##### CDs: State-dependent command mapping *No restriction.*

With *independent command* simulations,  $\mathcal{S}$  commands must be precompiled to  $\mathcal{T}$  commands which will work in any reachable  $\mathcal{T}$  state. This is a restriction placed by [10]. *Theory dependent* command mappings allow limited inspection of the  $\mathcal{T}$  state; this restriction allows the sequence of  $\mathcal{T}$  commands to be determined based only on the *theory* of the  $\mathcal{T}$  state: the values of all  $\mathcal{T}$  queries in the state. If two  $\mathcal{T}$  states answer all queries the same way, the same  $\mathcal{T}$  commands would be used in both to simulate an  $\mathcal{S}$  command. With this restriction, the monitor that transforms  $\mathcal{S}$  inputs into  $\mathcal{T}$  procedures need not be more privileged than users of the access control system, since queries are the user's only API to observe the state.

Finally, *state-dependent* command mappings can arbitrarily observe the state. This requires a monitor that is privileged enough to observe elements of the state that are not queriable, and two states that answer all queries identically may simulate commands differently depending on unobservable state.

#### Dimension CC: Command mapping complexity

Having considered the inputs available to the command mapping, we now consider the time complexity of this mapping. Note that this is measured as the increase in time as the *state* grows and thus is meaningless for independent command.

**CCc: Constant command mapping**  $\forall \psi \in \Psi^{\mathcal{S}}$ , the algorithm for  $\sigma_{\psi}(\gamma) = \sigma_{\Psi}(\psi, \gamma)$  has time complexity  $T(n) \in \mathcal{O}(1)$

**CCl: Linear command mapping**  $\forall \psi \in \Psi^{\mathcal{S}}$ , the algorithm for  $\sigma_{\psi}(\gamma) = \sigma_{\Psi}(\psi, \gamma)$  has time complexity  $T(n) \in \mathcal{O}(n)$

**CC $\infty$ : Unbounded command mapping** *No restriction.*

*Constant command* simulations do not allow more processing time for bigger states. Thus, the command mapping cannot loop over sets within the state. With *linear command*, the command mapping can take time linear in the size of the state, e.g., looping over sets in the state, but cannot contain double loops over sets, sort sets, etc. Finally, *unbounded command* simulations put no limit on the complexity of the command mapping (though we may expect it to have to be tractable, e.g., poly-time).

#### Dimension CS: Command mapping stuttering

Since the command mapping maps an  $\mathcal{S}$  state to a sequence of  $\mathcal{T}$  states, we may restrict the number of commands that can be used to simulate a single  $\mathcal{S}$  command.

**CS1: Lock-step**  $\forall \psi \in \Psi^{\mathcal{S}}, \gamma \in \Gamma^{\mathcal{T}} : |\sigma_{\Psi}(\psi, \gamma)| \leq 1$

**CS $c$ : Constant step**  $\exists c : \forall \psi \in \Psi^{\mathcal{S}}, \gamma \in \Gamma^{\mathcal{T}} : |\sigma_{\Psi}(\psi, \gamma)| \leq c$

**CS $\infty$ : Unbounded step** *No restriction.*

A *lock-step* simulation allows at most one  $\mathcal{T}$  command for each simulated  $\mathcal{S}$  command. This mitigates concurrency issues for multiuser systems, since the system does not pass through potentially inconsistent states between command executions. *Constant step* simulations allow multiple commands to be used, but only a number constant in the size of the state. Thus, multiple actions can be taken, but not, e.g., a command for each user in the system. Finally, *unbounded step* does not restrict how many  $\mathcal{T}$  commands can be executed per  $\mathcal{S}$  command.

#### Dimension CT: Trace structure

This class of properties enforces structural constraints on the traces of commands returned by the command mapping. This can address the potentially inconsistent states between start and end states in traces generated by the command mapping. Here, we present several examples of trace restrictions, using the notation *terminal*( $\gamma, \psi_1, \dots, \psi_j$ ) to denote the end state resulting from executing the sequence of commands  $\psi_1, \dots, \psi_j$ , starting from the state  $\gamma$ . Note that this dimension of properties is not totally ordered.




Fig. 3: A graphical representation of semantic lock-step

### CT1: Semantic lock-step

$$\begin{aligned}
& \forall \psi \in \Psi^S, \gamma^S \in \Gamma^S, \gamma^T \in \Gamma^T. ( \\
& \quad \exists \bar{\psi} = \langle \psi_1, \psi_2, \dots, \psi_m \rangle \in (\Psi^T)^*, i \in (1, m]. ( \\
& \quad \quad \sigma_\Psi(\psi, \gamma^T) = \bar{\psi} \wedge \\
& \quad \quad \forall j \in [1, i). (\gamma^S \sim \gamma^T \Rightarrow \\
& \quad \quad \quad \gamma^S \sim \text{terminal}(\gamma^T, \psi_1 \dots \psi_j)) \wedge \\
& \quad \quad \forall j \in [i, m]. (\gamma^S \sim \gamma^T \Rightarrow \\
& \quad \quad \quad \text{next}(\gamma^S, \psi) \sim \text{terminal}(\gamma^T, \psi_1 \dots \psi_j)))
\end{aligned}$$

First, a *semantic lock-step* simulation can appear to be lock-step (i.e., it does not enter any inconsistent states), because even though it is allowed to execute multiple  $\mathcal{T}$  commands to simulate a single  $\mathcal{S}$  command, only one of those commands is allowed to make correspondence-related changes. That is, consider the sequence of  $\mathcal{T}$  states constructed by executing the sequence of commands  $\sigma_\Psi(\psi^S, \gamma^T)$ . In semantic lock-step, all of these states must correspond to either the start state in  $\mathcal{S}$  or the end state in  $\mathcal{S}$ , and once the transition from start state to end state is made, the remaining states must all be equivalent to the end state. Thus, from the point of view of a user who can ask any combination of queries, the simulation appears to be lock-step. This restriction is depicted in Fig. 3.

### CTq: Query monotonic

$\forall \psi \in \Psi^S, \gamma \in \Gamma^T, q \in Q^T. \text{monotonic}(\psi, \gamma, q)$ , where:

$$\begin{aligned}
& \text{monotonic}(\psi, \gamma, q) \triangleq \exists \bar{\psi} = \langle \psi_1, \psi_2, \dots, \psi_m \rangle \in (\Psi^T)^*. ( \\
& \quad \sigma_\Psi(\psi, \gamma) = \bar{\psi} \wedge \\
& \quad \forall i \in (1, m). ( \\
& \quad \quad \text{terminal}(\gamma, \psi_1 \dots \psi_i) \vdash q \Rightarrow \\
& \quad \quad (\text{terminal}(\gamma, \psi_1 \dots \psi_{i-1}) \vdash q \vee \text{terminal}(\gamma, \psi_1 \dots \psi_m) \vdash q) \wedge \\
& \quad \quad \text{terminal}(\gamma, \psi_1 \dots \psi_i) \not\vdash q \Rightarrow \\
& \quad \quad (\text{terminal}(\gamma, \psi_1 \dots \psi_{i-1}) \not\vdash q \vee \text{terminal}(\gamma, \psi_1 \dots \psi_m) \not\vdash q))
\end{aligned}$$

Consider the start and end states of a trace in  $\mathcal{T}$ ,  $\gamma$  and  $\gamma'$ , respectively. Let  $Q^+$  be the set of queries that become true in  $\gamma'$  that were false in  $\gamma$ , and  $Q^-$  be the set of queries that become false in  $\gamma'$  that were true in  $\gamma$ . During the trace from  $\gamma$  to  $\gamma'$ , *query monotonicity* enforces that no queries are made true except  $Q^+$ , and no queries are made false except  $Q^-$ . Thus, from the point of view of a user who can ask only single queries, the simulation appears to be lock-step.

### CTa: Access monotonic

$\forall \psi \in \Psi^S, \gamma \in \Gamma^T, r \in \mathcal{R}^T. \text{monotonic}(\psi, \gamma, r)$

*Access monotonicity* is similar to query monotonicity but considering only authorization requests. Let  $\mathcal{R}^+$  be the set of requests that become allowed in  $\gamma'$  that were denied in  $\gamma$ , and  $\mathcal{R}^-$  be the set of requests that become denied in  $\gamma'$  that were allowed in  $\gamma$ . During the trace from  $\gamma$  to  $\gamma'$ , access

monotonicity enforces that no requests are granted except  $\mathcal{R}^+$ , and no requests are revoked except  $\mathcal{R}^-$ .

### CTs: Non-contaminating

$$\begin{aligned}
& \forall \psi \in \Psi^S, \gamma^T \in \Gamma^T. ( \\
& \quad \exists \bar{\psi} = \langle \psi_1, \psi_2, \dots, \psi_m \rangle \in (\Psi^T)^*. ( \\
& \quad \quad \sigma_\Psi(\psi, \gamma^T) = \bar{\psi} \wedge \\
& \quad \quad \forall \gamma_i^T \in \{ \gamma_i^T \mid \exists \psi_i \in \bar{\psi} : \gamma_i^T = \text{terminal}(\gamma^T, \psi_1 \dots \psi_i) \}. ( \\
& \quad \quad \quad \text{Allowed}(\gamma_i^T) \subseteq \text{Allowed}(\gamma^T) \vee \\
& \quad \quad \quad \text{Allowed}(\gamma_i^T) \subseteq \text{Allowed}(\text{terminal}(\gamma^T, \bar{\psi})))
\end{aligned}$$

The *non-contaminating* trace property ensures that no two accesses are allowed in the same state that are not both allowed in either the start or end state. This prevents, e.g., an intermediate state where a file can be accessed by two users simultaneously when simulating a command intended to *switch* which user can access the file. This definition uses the  $\text{Allowed}(\gamma)$  notation, indicating the set of all permissions  $p$  allowed in state  $\gamma$  (i.e., such that  $\gamma \vdash p$ ).

### Dimension CA: Actor preservation

Actor preservation properties restrict which users can be invoked in  $\mathcal{T}$  to handle  $\mathcal{S}$  commands. Here, we assume that  $\alpha(\psi)$  denotes the actor executing the command  $\psi$ . Note that this requires system support (e.g., the executing actor being an implicit argument passed to a command) in order for a simulation to be executable.

**CA $\mathcal{T}$ : Self-execution**  $\forall \psi^S \in \Psi^S, \gamma \in \Gamma^T, \forall \psi^T \in \sigma_\Psi(\psi^S, \gamma), \alpha(\psi^S) = \alpha(\psi^T)$

**CAa: Administration-preservation** Let  $A$  be the administrative subset of executing entities in the system.  $\forall \psi^S \in \Psi^S, \gamma \in \Gamma^T, \forall \psi^T \in \sigma_\Psi(\psi^S, \gamma), \alpha(\psi^T) \in A \Rightarrow \alpha(\psi^S) \in A$

*Self-execution* says that any command in  $\mathcal{S}$  executed by any user  $u$  must be mapped to a sequence of commands in  $\mathcal{T}$ , all of which are executed by  $u$ . *Administration-preservation* prevents the invocation of administrators in  $\mathcal{T}$  where they were not needed in  $\mathcal{S}$ . In an administration-preserving simulation, any command in  $\mathcal{S}$  executed by a non-administrative user is mapped to a sequence of commands in  $\mathcal{T}$ , none of which is executed by an administrator. Other forms of actor preservation, as well as defining the set of administrators, are application-specific.

### D. Query decider properties

We defer the bulk of the technical discussion of the query decider restrictions to Appendix A, as they are largely similar to the command mapping restrictions. Query decider dependence (QD), like command mapping dependence (CD), restricts the information that the query decider can consider about a  $\mathcal{T}$  state when deciding the truth value of an  $\mathcal{S}$  query in that state. Query decider complexity (QC) restricts the runtime of the routine.

Query preservation (QP) indicates which queries need to stay the same as they are mapped from system  $\mathcal{S}$  to system  $\mathcal{T}$ . A particular application may require any given set of queries to be preserved; the most common property in this dimension



is *authorization preservation*, which enforces that the query decider maps each  $\mathcal{S}$  request to the value of the identical request in the  $\mathcal{T}$  state. This can be seen as ensuring that  $\mathcal{T}$  is using its model “as intended” (i.e., forcing it to answer simulated requests as it would its own native requests).

#### E. Reachability

##### Dimension R: Reachability

The last dimension of properties we consider ties the mappings together to ensure the simulation is indeed what one could consider a simulation in the classic sense. A state correspondence, query decider, and command mapping do not automatically define a simulation without reachability constraints. Here, we define forward and bidirectional reachability, two variants of this type of constraint (note that these properties are presented in *increasing* strictness since the latter builds upon the former).

##### $R \rightarrow$ : Forward reachability

$$\begin{aligned} \forall \gamma_0^S, \gamma_1^S \in \Gamma^S, \gamma_0^T \in \Gamma^T. ( \\ \gamma_0^S \sim \gamma_0^T \wedge \gamma_0^S \mapsto \gamma_1^S \Rightarrow \exists \gamma_1^T \in \Gamma^T. ( \\ \gamma_0^T \xrightarrow{*} \gamma_1^T \wedge \gamma_1^S \sim \gamma_1^T)) \end{aligned}$$

##### $R \leftrightarrow$ : Bidirectional reachability *Forward reachability, and:*

$$\begin{aligned} \forall \gamma_0^S \in \Gamma^S, \gamma_0^T, \gamma_1^T \in \Gamma^T. ( \\ \gamma_0^S \sim \gamma_0^T \wedge \gamma_0^T \mapsto \gamma_1^T \Rightarrow \exists \gamma_1^S \in \Gamma^S. ( \\ \gamma_0^S \xrightarrow{*} \gamma_1^S \wedge \gamma_1^S \sim \gamma_1^T)) \end{aligned}$$

In *forward reachability*, any transition made in  $\mathcal{S}$  must be possible in  $\mathcal{T}$ . If  $\gamma_0^S$  corresponds to  $\gamma_0^T$ , and  $\gamma_1^S$  can be reached from  $\gamma_0^S$  via the commands of  $\mathcal{S}$ , then  $\gamma_1^S$  must correspond to a state  $\gamma_1^T$  in  $\mathcal{T}$  that is reachable from  $\gamma_0^T$ . The notion of state correspondence is determined by the property chosen in dimension SC.

*Bidirectional reachability* (or bi-reachability), also requires that  $\mathcal{T}$  cannot enter a state that does not correspond to a reachable state in  $\mathcal{S}$ . If  $\gamma_0^S$  corresponds to  $\gamma_0^T$ , and  $\gamma_1^T$  is reachable from  $\gamma_0^T$  by executing a command, then there must exist an  $\mathcal{S}$  state  $\gamma_1^S$  that corresponds to  $\gamma_1^T$  and that is reachable from  $\gamma_0^S$  by executing one or more commands. This process may make use of multiple steps, since the procedure for finding the corresponding  $\mathcal{S}$  states does not need to be constructed, these states must simply exist. The operational advantage of enforcing  $R \leftrightarrow$  is that, even if the simulating system’s native operations are exposed to users, the system can never enter a state that does not have an equivalent in the simulated system.

## V. POSITIONING EXISTING SIMULATIONS

As mentioned in Section IV-A, no set of properties can be proven to describe all conceivable simulations. In this section, we support the set of properties defined in this work by showing that it can *precisely* describe the wide range of existing expressiveness simulations.

#### A. Expressiveness using Simulation Properties

We will now draw the formal distinction between a simulation and expressiveness. Here, we use  $\mathcal{T} \overset{\mathcal{X}}{\text{sim}} \mathcal{S}$  to denote, “ $\mathcal{T}$  can admit a simulation of type  $\mathcal{X}$  of  $\mathcal{S}$ ,” and  $\mathcal{S} \leq^{\mathcal{X}} \mathcal{T}$  to denote, “ $\mathcal{T}$  is at least as expressive as  $\mathcal{S}$  with respect to simulations of type  $\mathcal{X}$ .”

While previous work considers the expressiveness result to be equivalent to a simulation (i.e.,  $\mathcal{T} \overset{\mathcal{X}}{\text{sim}} \mathcal{S} \equiv \mathcal{S} \leq^{\mathcal{X}} \mathcal{T}$ ), expressiveness in a practical sense is subject to a subtle distinction. Since we mean for expressiveness to be *implementable* (i.e., if  $\mathcal{T}$  is as expressive as  $\mathcal{S}$ , then  $\mathcal{T}$  can be used in place of  $\mathcal{S}$ ), expressiveness within the domain of simulation properties should mean the following: if  $\mathcal{T}$  is as expressive as  $\mathcal{S}$ , then  $\mathcal{T}$  can simulate any system that  $\mathcal{S}$  can simulate. Thus, we define *expressiveness* in the context of a set of simulation properties.

**Definition 6 (Expressiveness)** *Given access control systems  $\mathcal{S}$  and  $\mathcal{T}$  and a set of simulation properties  $\mathcal{P}$ , we say that  $\mathcal{T}$  is at least as expressive as  $\mathcal{S}$  with respect to  $\mathcal{P}$  (denoted  $\mathcal{S} \leq^{\mathcal{P}} \mathcal{T}$ ) to mean that, for every system  $\mathcal{U}$ , if  $\mathcal{S}$  can simulate  $\mathcal{U}$  while enforcing  $\mathcal{P}$ , then  $\mathcal{T}$  can simulate  $\mathcal{U}$  while enforcing  $\mathcal{P}$  ( $\forall \mathcal{U} : \mathcal{S} \overset{\mathcal{P}}{\text{sim}} \mathcal{U} \Rightarrow \mathcal{T} \overset{\mathcal{P}}{\text{sim}} \mathcal{U}$ ).*  $\diamond$

We first point out that this definition of expressiveness is strictly more general than the more traditional (often implied) notion. Since  $\mathcal{S}$  can trivially simulate itself,  $\mathcal{S} \leq^{\mathcal{X}} \mathcal{T}$  implies  $\mathcal{T} \overset{\mathcal{X}}{\text{sim}} \mathcal{S}$ . The additional generalization can be viewed from a formal standpoint as dropping the (incorrect) assumption that all types of simulation are transitive (i.e., that  $\mathcal{T} \overset{\mathcal{X}}{\text{sim}} \mathcal{S}$  and  $\mathcal{S} \overset{\mathcal{X}}{\text{sim}} \mathcal{U}$  imply  $\mathcal{T} \overset{\mathcal{X}}{\text{sim}} \mathcal{U}$ ). For instance, assume that  $\mathcal{T}$  can simulate  $\mathcal{S}$  and  $\mathcal{S}$  can simulate  $\mathcal{U}$ , each with a quadratic increase in state storage. While  $\mathcal{T}$  may be able to simulate  $\mathcal{U}$ , this simulation may require greater than quadratic storage.

From a more intuitive standpoint, we point out that, except in the case of custom-built access control solutions, any deployment is a simulation of a workload (i.e., ideal operation) using an existing system. That is, unless  $\mathcal{S}$  is custom-made to exactly satisfy the desired workload, replacing it with  $\mathcal{T}$  is not a matter of whether  $\mathcal{T}$  can simulate  $\mathcal{S}$ , but whether  $\mathcal{T}$  can admit an equally good simulation of the (perhaps not formally specified) workload that  $\mathcal{S}$  is known to simulate. This concept is discussed by Kane and Browne [27], who point out that an access control implementation is often only an approximation of the desired policy. In particular, as policy languages get more complex, deployments often make use of approximations that are easier to analyze and more efficient to enforce than the overly-expressive policy language.

#### B. Decomposing Expressiveness Simulations to Properties

In order to use the set of expressiveness simulation properties detailed in Section IV to systematically compare previously proposed notions of simulation, we present our formal way of stating that a notion of simulation and a set of simulation properties are equivalent. We call this correspondence *simulation decomposition*: when a notion of simulation  $\mathcal{X}$  can be

decomposed to a set of simulation properties  $\mathcal{P}$ , then analyses using  $\mathcal{X}$  and  $\mathcal{P}$  yield equivalent expressiveness results.

**Definition 7 (Simulation Decomposition)** *Given a notion of access control simulation  $\mathcal{X}$  and a set of simulation properties  $\mathcal{P}$ ,  $\mathcal{X}$  can be decomposed to  $\mathcal{P}$  (denoted  $\mathcal{X} \doteq \mathcal{P}$ ) if and only if, for all systems  $\mathcal{S}$  and  $\mathcal{T}$ ,  $\mathcal{T} \stackrel{\mathcal{X}}{\text{sim}} \mathcal{S} \iff \mathcal{S} \leq^{\mathcal{P}} \mathcal{T}$ . That is,  $\mathcal{T}$  admits an  $\mathcal{X}$  simulation of  $\mathcal{S}$  if and only if  $\mathcal{T}$  is at least as expressive as  $\mathcal{S}$  with respect to properties  $\mathcal{P}$ .*  $\diamond$

Recall from Definition 6 that  $\mathcal{S} \leq^{\mathcal{P}} \mathcal{T}$  says that any system that can be simulated by  $\mathcal{S}$  while preserving properties  $\mathcal{P}$  can also be simulated by  $\mathcal{T}$  while preserving  $\mathcal{P}$ . In light of this, we will position an existing notion of simulation,  $\mathcal{X}$ , within the lattice formed by our simulation properties (i.e., prove  $\mathcal{X} \doteq \mathcal{P}$ ) by proving the following for the set of properties  $\mathcal{P}$ :

- 1) (Only-if direction)  $\mathcal{T} \stackrel{\mathcal{X}}{\text{sim}} \mathcal{S} \wedge \mathcal{S} \stackrel{\mathcal{P}}{\text{sim}} \mathcal{U} \Rightarrow \mathcal{T} \stackrel{\mathcal{P}}{\text{sim}} \mathcal{U}$
- 2) (If direction)  $\mathcal{S} \leq^{\mathcal{P}} \mathcal{T} \Rightarrow \mathcal{T} \stackrel{\mathcal{X}}{\text{sim}} \mathcal{S}$

We give an example of such a proof in the following section.

### C. Example Decomposition

To demonstrate how simulation decomposition proofs are written, we now consider the Ammann-Lipton-Sandhu simulation [6]. The ALS simulation considers access control states as graphs: sets of primitive objects are node types, and sets of relations are edge types. The set of node types and edge types in the states of system  $\mathcal{S}$  are denoted  $NT(\mathcal{S})$  and  $ET(\mathcal{S})$ , respectively. The ALS state correspondence is then defined as follows (reworded slightly from [6]).

**Definition 8** *A state in system  $\mathcal{S}$ , a simulated system, and a state in system  $\mathcal{T}$ , a simulating system, correspond iff the graph defining the state in  $\mathcal{S}$  is identical to the subgraph obtained by taking the state in  $\mathcal{T}$  and discarding all nodes (edges) not in  $NT(\mathcal{S})$  ( $ET(\mathcal{S})$ ).*  $\diamond$

The ALS simulation is defined with respect to this state correspondence.

**Definition 9** *Under the definition of correspondence in Definition 8, system  $\mathcal{T}$  simulates system  $\mathcal{S}$  iff the following conditions hold:*

- 1) *If system  $\mathcal{S}$  can reach a given state, system  $\mathcal{T}$  can reach a corresponding state.*
- 2) *If system  $\mathcal{T}$  can reach a given state, system  $\mathcal{S}$  can reach a corresponding state.*

We will now demonstrate the two-step simulation decomposition proof technique described in Section V-B for the ALS simulation. For the purposes of this proof, let the set of simulation properties  $\mathcal{P} = \{\text{SCs}, \text{QPa}, \text{R}\leftrightarrow\}$ . Recall that SCs is structure state correspondence, which says that the simulating state must include all of the unaltered sets from the simulated state; QPa is authorization preservation, which says that each authorization request must be mapped identically from simulated to simulating system (and thus the simulating system must support the same set of requests as the simulated system); and  $\text{R}\leftrightarrow$  is bireachability, which says that the simulating system

can reach a state which corresponds to each reachable simulated state, and cannot reach a state which does not correspond to a reachable state in the simulated system.

We will demonstrate the two steps of the proof technique by proving two requisite lemmas. First, step 1 (only-if direction):

**Lemma 1** *Given access control systems  $\mathcal{S}$ ,  $\mathcal{T}$ , and  $\mathcal{U}$ ,*

$$\mathcal{T} \stackrel{\text{ALS}}{\text{sim}} \mathcal{S} \wedge \mathcal{S} \stackrel{\mathcal{P}}{\text{sim}} \mathcal{U} \Rightarrow \mathcal{T} \stackrel{\mathcal{P}}{\text{sim}} \mathcal{U}$$

*That is, if  $\mathcal{T}$  admits an ALS simulation of  $\mathcal{S}$ , and  $\mathcal{S}$  admits a simulation of  $\mathcal{U}$  with properties  $\{\text{SCs}, \text{QPa}, \text{R}\leftrightarrow\}$ , then  $\mathcal{T}$  admits a simulation of  $\mathcal{U}$  with properties  $\{\text{SCs}, \text{QPa}, \text{R}\leftrightarrow\}$ .*

*Proof:* To prove this lemma, we let  $\mathcal{S}$ ,  $\mathcal{T}$ , and  $\mathcal{U}$  be access control systems such that  $\mathcal{T} \stackrel{\text{ALS}}{\text{sim}} \mathcal{S}$  and  $\mathcal{S} \stackrel{\mathcal{P}}{\text{sim}} \mathcal{U}$  but are otherwise arbitrary, and we show that  $\mathcal{T} \stackrel{\mathcal{P}}{\text{sim}} \mathcal{U}$ .

Choose an arbitrary state  $\gamma_0^{\mathcal{U}} \in \Gamma^{\mathcal{U}}$  and command  $\psi^{\mathcal{U}} \in \Psi^{\mathcal{U}}$ , and let  $\text{next}(\gamma_0^{\mathcal{U}}, \psi^{\mathcal{U}}) = \gamma_1^{\mathcal{U}}$ . Let  $\gamma_0^{\mathcal{S}} \in \Gamma^{\mathcal{S}}$  such that  $\gamma_0^{\mathcal{U}} \stackrel{s}{\sim} \gamma_0^{\mathcal{S}}$ . Since  $\mathcal{S} \stackrel{\mathcal{P}}{\text{sim}} \mathcal{U}$ ,

$$\exists \gamma_1^{\mathcal{S}} \in \Gamma^{\mathcal{S}}. (\text{terminal}(\gamma_0^{\mathcal{S}}, \sigma_{\Psi}(\psi^{\mathcal{U}}, \gamma_0^{\mathcal{S}})) = \gamma_1^{\mathcal{S}} \wedge \gamma_1^{\mathcal{U}} \stackrel{s}{\sim} \gamma_1^{\mathcal{S}})$$

Let  $\gamma_0^{\mathcal{T}} \in \Gamma^{\mathcal{T}}$  such that  $\gamma_0^{\mathcal{S}} \stackrel{s}{\sim} \gamma_0^{\mathcal{T}}$ . Since  $\mathcal{T} \stackrel{\text{ALS}}{\text{sim}} \mathcal{S}$ ,

$$\exists \gamma_1^{\mathcal{T}} \in \Gamma^{\mathcal{T}}. (\gamma_0^{\mathcal{T}} \mapsto^* \gamma_1^{\mathcal{T}} \wedge \gamma_1^{\mathcal{S}} \stackrel{s}{\sim} \gamma_1^{\mathcal{T}})$$

Thus, there exists a sequence of  $\mathcal{T}$  commands  $\Psi_0^{\mathcal{T}}$  such that  $\text{terminal}(\gamma_0^{\mathcal{T}}, \Psi_0^{\mathcal{T}}) = \gamma_1^{\mathcal{T}}$ . Define  $\sigma_{\Psi} : \Psi^{\mathcal{U}} \times \Gamma^{\mathcal{T}} \rightarrow (\Psi^{\mathcal{T}})^*$  such that it returns  $\Psi_0^{\mathcal{T}}$  for  $\gamma_0^{\mathcal{T}}, \psi^{\mathcal{U}}$ .

Then, given  $\gamma_0^{\mathcal{U}}, \gamma_1^{\mathcal{U}} \in \Gamma^{\mathcal{U}}, \gamma_0^{\mathcal{T}} \in \Gamma^{\mathcal{T}}, \psi^{\mathcal{U}} \in \Psi^{\mathcal{U}}$  such that  $\text{next}(\gamma_0^{\mathcal{U}}, \psi^{\mathcal{U}}) = \gamma_1^{\mathcal{U}}$ , and  $\gamma_0^{\mathcal{U}} \stackrel{s}{\sim} \gamma_0^{\mathcal{T}}$ ,

$$\exists \gamma_1^{\mathcal{T}} \in \Gamma^{\mathcal{T}}. (\text{terminal}(\gamma_0^{\mathcal{T}}, \sigma_{\Psi}(\psi, \gamma_0^{\mathcal{T}})) = \gamma_1^{\mathcal{T}} \wedge \gamma_1^{\mathcal{S}} \stackrel{s}{\sim} \gamma_1^{\mathcal{T}})$$

Hence,  $\mathcal{T} \stackrel{\{\text{SCs}, \text{R}\leftrightarrow\}}{\text{sim}} \mathcal{U}$ . Next, we show QPa.

Choose some arbitrary request  $r_0^{\mathcal{U}} \in \mathcal{R}^{\mathcal{U}}$  and state  $\gamma_0^{\mathcal{T}} \in \Gamma^{\mathcal{T}}$ . Since  $\mathcal{S} \stackrel{\mathcal{P}}{\text{sim}} \mathcal{U}$ ,

$$\forall r^{\mathcal{U}} \in \mathcal{R}^{\mathcal{U}}, \gamma^{\mathcal{S}} \in \Gamma^{\mathcal{S}}, \sigma_Q(r^{\mathcal{U}}, \gamma^{\mathcal{S}}) = \gamma^{\mathcal{S}} \vdash r^{\mathcal{U}}$$

Thus, we know that  $\mathcal{S}$  supports all  $\mathcal{U}$  requests, and corresponding  $\mathcal{S}$  and  $\mathcal{U}$  states will answer  $\mathcal{U}$  requests identically. Therefore,  $r_0^{\mathcal{U}} \in \mathcal{R}^{\mathcal{S}}$ . Since  $\mathcal{T} \stackrel{\text{ALS}}{\text{sim}} \mathcal{S}$ ,

$$\forall r^{\mathcal{S}} \in \mathcal{R}^{\mathcal{S}}, \gamma^{\mathcal{T}} \in \Gamma^{\mathcal{T}}, \sigma_Q(r^{\mathcal{S}}, \gamma^{\mathcal{T}}) = \gamma^{\mathcal{T}} \vdash r^{\mathcal{S}}$$

Thus,  $\sigma_Q(r_0^{\mathcal{U}}, \gamma_0^{\mathcal{T}}) = \gamma_0^{\mathcal{T}} \vdash r_0^{\mathcal{U}}$ .

Hence,  $\mathcal{T} \stackrel{\{\text{SCs}, \text{QPa}, \text{R}\leftrightarrow\}}{\text{sim}} \mathcal{U}$ . Next, we show  $\text{R}\leftrightarrow$ .

Choose some arbitrary states  $\gamma_0^{\mathcal{T}}, \gamma_1^{\mathcal{T}} \in \Gamma^{\mathcal{T}}$  such that  $\gamma_0^{\mathcal{T}} \mapsto^* \gamma_1^{\mathcal{T}}$ . Let  $\gamma_0^{\mathcal{S}} \in \Gamma^{\mathcal{S}}$  such that  $\gamma_0^{\mathcal{T}} \stackrel{s}{\sim} \gamma_0^{\mathcal{S}}$ . Since  $\mathcal{T} \stackrel{\text{ALS}}{\text{sim}} \mathcal{S}$ ,


$$\exists \gamma_1^{\mathcal{S}}. (\gamma_0^{\mathcal{S}} \mapsto^* \gamma_1^{\mathcal{S}} \wedge \gamma_1^{\mathcal{S}} \stackrel{s}{\sim} \gamma_1^{\mathcal{T}})$$

Let  $\gamma_0^{\mathcal{U}} \in \Gamma^{\mathcal{U}}$  such that  $\gamma_0^{\mathcal{U}} \stackrel{s}{\sim} \gamma_0^{\mathcal{S}}$ . Since  $\mathcal{S} \stackrel{\mathcal{P}}{\text{sim}} \mathcal{U}$ ,


$$\exists \gamma_1^{\mathcal{U}}. (\gamma_0^{\mathcal{U}} \mapsto^* \gamma_1^{\mathcal{U}} \wedge \gamma_1^{\mathcal{U}} \stackrel{s}{\sim} \gamma_1^{\mathcal{S}})$$

Simulation	Decomposition
ALS	SCs QPa R $\leftrightarrow$
CDMw	SCa QPa CDi R $\rightarrow$
CDMs	SCa QPa CDi CS1 R $\rightarrow$
Ganta	SCa QPa CTs R $\leftrightarrow$
HMG+	SCq QDt R $\rightarrow$
HMG+a	QPa
HMG+s	CTa
HMG+p	CAa
SMG	SCa R $\rightarrow$
TL-SMR	SCq QD1 R $\leftrightarrow$

(a) Decompositions of surveyed simulations



(b) Taxonomy of simulations



(c) Partial lattice of simulations

Fig. 4

Thus, given  $\gamma_0^T, \gamma_1^T \in \Gamma^T, \gamma_0^U \in \Gamma^U$  such that  $\gamma_0^T \mapsto \gamma_1^T$  and  $\gamma_0^U \stackrel{s}{\sim} \gamma_0^T$ ,

$$\exists \gamma_1^U \in \Gamma^U. (\gamma_0^U \mapsto \gamma_1^U \wedge \gamma_1^U \stackrel{s}{\sim} \gamma_1^T)$$

Hence,  $\mathcal{T} \stackrel{sim}{\sim} \mathcal{U}$ .

Next, we demonstrate step 2 (if direction):

**Lemma 2** Given access control systems  $\mathcal{S}$  and  $\mathcal{T}$  and simulation properties  $\mathcal{P} = \{\text{SCs}, \text{QPa}, \text{R}\leftrightarrow\}$ ,  $\mathcal{S} \leq^{\mathcal{P}} \mathcal{T} \Rightarrow \mathcal{T} \stackrel{sim}{\sim}_{ALS} \mathcal{S}$ . That is, if  $\mathcal{T}$  is at least as expressive as  $\mathcal{S}$  with respect to properties  $\mathcal{P}$ , then  $\mathcal{T}$  admits an ALS simulation of  $\mathcal{S}$ .

*Proof:* To prove this lemma, we let  $\mathcal{S}$  and  $\mathcal{T}$  be arbitrary access control systems such that  $\mathcal{S} \leq^{\mathcal{P}} \mathcal{T}$ , and we show that  $\mathcal{T} \stackrel{sim}{\sim}_{ALS} \mathcal{S}$ .

Since  $\mathcal{S} \leq^{\mathcal{P}} \mathcal{T}$ , for any access control system  $\mathcal{U}$ , if  $\mathcal{S} \stackrel{sim}{\sim}_{\mathcal{P}} \mathcal{U}$ , then  $\mathcal{T} \stackrel{sim}{\sim}_{\mathcal{P}} \mathcal{U}$ .

Since  $\mathcal{S}$  can trivially simulate itself,  $\mathcal{S} \stackrel{sim}{\sim}_{\mathcal{P}} \mathcal{S}$ , and thus  $\mathcal{T} \stackrel{sim}{\sim}_{\mathcal{P}} \mathcal{S}$ .

Thus, given  $\gamma_0^S, \gamma_1^S \in \Gamma^S, \gamma_0^T \in \Gamma^T$ , by forward reachability, if  $\gamma_0^S \stackrel{s}{\sim} \gamma_0^T$  and  $\gamma_0^S \mapsto \gamma_1^S$ , then

$$\exists \gamma_1^T. (\gamma_0^T \mapsto \gamma_1^T \wedge \gamma_1^S \stackrel{s}{\sim} \gamma_1^T)$$

Since SCs and QPa satisfy the ALS definition of state correspondence, this means we have satisfied the first property of the ALS simulation.

1) If  $\mathcal{S}$  can reach a given state,  $\mathcal{T}$  can reach a corresponding state.

And by bidirectional reachability, given  $\gamma_0^S \in \Gamma^S, \gamma_0^T, \gamma_1^T \in \Gamma^T$ , if  $\gamma_0^S \stackrel{s}{\sim} \gamma_0^T$  and  $\gamma_0^T \mapsto \gamma_1^T$ , then

$$\exists \gamma_1^S. (\gamma_0^S \mapsto \gamma_1^S \wedge \gamma_1^S \stackrel{s}{\sim} \gamma_1^T)$$

And therefore, we have satisfied the second property of the ALS simulation:

2) If  $\mathcal{T}$  can reach a given state,  $\mathcal{S}$  can reach a corresponding state.

These properties satisfy the definition for ALS simulation, and hence  $\mathcal{T}$  admits an ALS simulation of  $\mathcal{S}$  ( $\mathcal{T} \stackrel{sim}{\sim}_{ALS} \mathcal{S}$ ). ■

Therefore, we have proved the decomposition of the ALS simulation:

**Theorem 3**  $ALS \equiv \{\text{SCs}, \text{QPa}, \text{R}\leftrightarrow\}$ ; that is, the ALS simulation decomposes to structure correspondence, authorization preservation, and bidirectional reachability.

*Proof:* By Lemma 1, if  $\mathcal{T} \stackrel{sim}{\sim}_{ALS} \mathcal{S}$ , then  $\mathcal{S} \leq^{\mathcal{P}} \mathcal{T}$ . By Lemma 2, if  $\mathcal{S} \leq^{\mathcal{P}} \mathcal{T}$ , then  $\mathcal{T} \stackrel{sim}{\sim}_{ALS} \mathcal{S}$ . Thus,  $\mathcal{S} \leq^{\mathcal{P}} \mathcal{T}$  if and only if  $\mathcal{T} \stackrel{sim}{\sim}_{ALS} \mathcal{S}$ , and thus the ALS simulation decomposes to  $\{\text{SCs}, \text{QPa}, \text{R}\leftrightarrow\}$ . ■

All other decomposition proofs can be found in Appendix B.

#### D. Results

Now, we present the results of decomposing the simulations from the series of previous works discussed in Section II into sets of simulation properties from Section IV. First, a chart of our results is shown in Fig. 4a, which states the decomposition of the SMG simulation [2]–[4], [7]–[9], the Ganta simulation [5], the ALS simulation [1], [6], the CDM weak and strong simulations [10], the TL state-matching reduction [12], [13], and HMG+ parameterized expressiveness (along with several parameterized expressiveness properties) [14]. Properties are omitted if they are not explicitly required by the simulation's definition but are implied by other, explicit properties (e.g., CDMs decomposes to a set including CDi, which also implies CCc). Section VI-A discusses which properties imply others.

In Fig. 4b, we arrange this data as a taxonomy, with each split representing a dimension, with weaker properties positioned to the left and stronger properties to the right. We split first on the state correspondence, which is perhaps the biggest difference among the surveyed simulations. This separates simulations that preserve only the answers to authorization requests (SCa) from those that preserve all queries (SCq) and those that preserve full state structure (SCs). We note that the ALS simulation is alone in its decomposition including SCs; all other surveyed simulations allowed the simulating system to store information in a different organization than the simulated system, so long as the required querable information (requests or queries) can be recovered. We also note that the predominant difference between the SMG simulation and the CDM simulations is the command dependence: in SMG, a command can be mapped completely differently if it is to be executed in different states, while in CDM, each command must be mapped without knowing the state in which it will be executed. The Ganta simulation is unique in enforcing

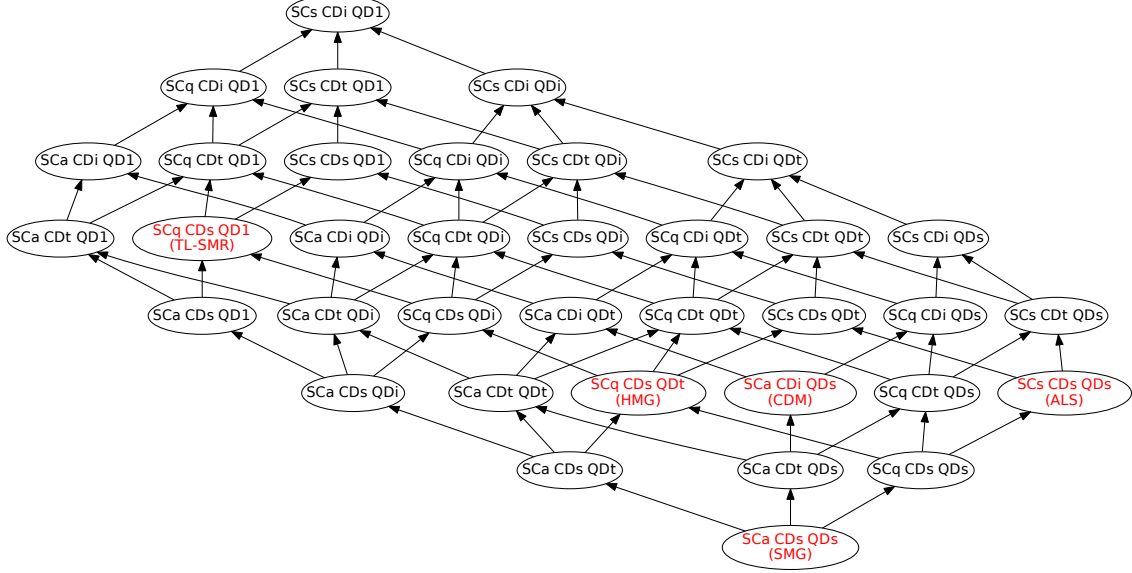


Fig. 5: Lattice of state correspondence, command dependence, and query dependence with positioned surveyed simulations

the non-contamination trace restriction. HMG+ and TL-SMR use the same state correspondence, but HMG+ enforces a more lax query dependence and does not require bireachability. Simulations that are positioned farther apart are the most dissimilar. Most starkly different are SMG and ALS, positioned far left and far right, which share no simulation properties except in dimensions in which both enforce only minimum properties, despite their similar publication times.

In Fig. 4c, we position the surveyed simulations within a lattice. Higher simulations decompose to more strict properties, and an arrow from simulation  $\mathcal{X}$  to simulation  $\mathcal{X}'$  indicates that  $\mathcal{X}'$  decomposes to strictly stronger properties than  $\mathcal{X}$ . Here we can see that the SMG simulation is strictly weakest, which supports previous claims to this effect [5], [13]. Several orthogonal directions were taken in defining other simulations to enforce stronger properties. The CDM simulations, as noted above, restrict the command dependence. The Ganta simulation requires non-contamination and bireachability. The TL state-matching reduction and HMG+ parameterized expressiveness consider queries, and thus strengthen the state correspondence. The ALS simulation enforces an even more strict state correspondence, requiring the structure of a simulated system's state to be preserved in the simulating system. Interestingly, we note that while all are stronger than SMG, most pairs are incomparable due to being stronger in orthogonal ways. In particular, while TL-SMR is considered to be a relatively strong notion of simulation, this is not substantiated by the lattice, which shows TL-SMR to be stronger than HMG+ and SMG, but incomparable to the CDM, ALS, and Ganta simulations.

Figure 5 presents a lattice of state correspondence, command dependence, and query dependence, with the surveyed simulations positioned within it (in this space, the Ganta simulation

is at the same point as the SMG simulation). This fig. makes evident the wide range of points between existing simulations that have not been explored. In this fig., we omit several dimensions for readability, namely reachability (which further separates Ganta, ALS, and TL-SMR from SMG, CDM and HMG+) and stuttering (which would break CDM into its weak and strong counterparts). Perhaps the most interesting points to explore within this lattice are those that exist between two surveyed simulations. For example,  $\{\text{SCq}, \text{CDs}, \text{QDs}\}$  adds to SMG the preservation of queries beyond requests, but stops short of HMG+ by not restricting the query decider to consider only the theory of the state while mapping queries. Similarly,  $\{\text{SCa}, \text{CDt}, \text{QDs}\}$  takes away some of SMG's freedom to inspect the state mapping commands, but rather than go all the way to the independent command mapping of CDM, it still allows it to inspect the state's responses to queries. We also point out  $\{\text{SCq}, \text{CDs}, \text{QDi}\}$ , which differs from HMG+ by enforcing query decider independence (mapping queries cannot consider the state or theory), but can map each simulated query to a boolean expression over simulating queries.

## VI. SELECTING NEW SETS OF PROPERTIES

In Section V, we positioned the simulations used in previous works within a comparative lattice, allowing them to be formally compared for the first time. In this section, we enable a second use of our lattice of expressiveness simulation properties: crafting new notions of expressiveness by choosing the properties that most closely correspond to the scenario in which an access control system will be deployed. We first discuss interactions between dimensions; this discussion should act as a warning against choosing individual properties in isolation. We then interpret the impact each identified dimension

has on the simulation, and identify properties of a deployment scenario that may dictate particular choices in each dimension. Finally, we discuss the potential impact these techniques could have on future expressiveness analysis.

#### A. Interactions Between Dimensions

We noted in Section V that some simulations decompose to sets of properties that include *implied* properties, or properties that are redundant given the others in the set. For instance, command independence (CDi) implies constant-time command mapping (CCc); if the command mapping does not depend on the state, then its procedure must be constant-time in the size of the state. Further, CCc implies constant step (CSc), since a constant-time procedure must have constant-size output.

An additional type of interaction is between basic properties and those properties whose definition relies on the basic properties in the abstract. For example, the definition of forward reachability ( $R \rightarrow$ ) refers to sequences of commands output by  $\sigma^\Psi$ , the length of which may be limited by command mapping stuttering (CS). Further, the definitions of both reachability properties (R) and trace structure properties (CT) refer to corresponding states. Here, the details of what makes states correspond is left to the state correspondence structure (SC).

These dependencies show that the proof of a property in one dimension may rely on the properties chosen in another. Thus, e.g., changing to a stronger state correspondence requires re-proving a simulation's results for reachability and trace structure, since these are dependent on state correspondence.

Several property dimensions are defined over the size of the simulated state: command mapping complexity (CC), command mapping stuttering (CS), and query decider time-complexity (QC). Thus, these dimensions can be altered with respect to the original, simulated state by the state storage size (SS). For example, enforcing polynomial storage (SSp) and linear-time command mapping (CCl) will guarantee a command mapping that is linear-time with respect to the *simulating state*, which is a polynomial expansion over the original *simulated state*.

#### B. Interpreting the Dimensions

We now discuss the practical impacts each identified dimensions, and what types of environments may cause one to prefer a particular property in these dimensions over others.

**SC: State correspondence structure** allows one to change what needs to be preserved about the state during a simulation. If the deployment scenario in question assumes only that the simulation allows the proper authorization requests, SCa should suffice. For scenarios that require the access control system to support (and provide correct answers to) additional queries such as, “*Is user  $u$  a member of role  $r$ ?*”, SCq is more appropriate. Finally, in scenarios that make use of additional code that has access to (and assumes a particular arrangement of) the access control system's internal data structures, SCs is the best choice.

**SS: State storage** limits the size of the simulated state with respect to the original state (i.e., the state of the system being simulated). This can be restricted for several reasons. The most obvious is storage space: if trusted storage for representing

access control state is limited, we may restrict the simulation from mapping states in a way that increases storage by more than a linear factor (SSl) or a polynomial factor (SSp). However, the more interesting reason comes from an interaction described in Section VI-A. Since other dimensions place restrictions (e.g., on the number of commands executed) based on the size of the simulating state, we may restrict the state expansion to linear (SSl) in order, e.g., to restrict the command mapping procedure to be linear-time in the size of the *original, simulated state*. If state storage is polynomial (SSp), then even if we enforce a command mapping that is linear in the simulating state (CCl), this only restricts it to being polynomial-time with respect to the simulated state. Thus, even when trusted storage space is unbounded in the deployment scenario, one may desire to limit state size to limit later iteration over this state.

**CD: Command mapping dependence** allows one to require that the command mapping be computable without full knowledge and inspection of the state in which a command will be executed in. Independent command (CDi) requires that each command is mapped independent of the state, and is useful in deployment scenarios in which the agent calculating the simulating commands is completely unprivileged, and cannot inspect the state. It is also useful when commands must be precompiled, thus adding no computation at runtime beyond that of the simulating commands themselves. Theory-dependent command mapping (CDt) allows the command mapping to inspect the *theory* of the state (i.e., the answers to all queries). This property is useful in deployment scenarios in which the simulation agent is no more privileged than normal users—calculating the mapped commands requires only information available by asking queries. Finally, state-dependent command mapping (CDs) allows the command mapping to arbitrarily inspect the state, requiring a powerful simulation agent.

**CC: Command mapping complexity** restricts the time-complexity of the command mapping with respect to the size of the simulating state. Constant command mapping (CCc) can restrict the command mapping from taking any longer for larger states, and is thus appropriate when states can be large but mapping commands must always remain fast. Linear command mapping (CCl) prevents expensive nested loops over access control state as well as operations such as sorting, while still allowing more processing for larger states.

**CS: Command stuttering** restricts the number of simulating commands executed for each simulated command. Lock-step (CS1) simulations must execute no more than one simulating command per simulated command, and thus ensure there is no intermediate state exposed to users. In deployment scenarios without the ability to force atomic execution of a sequence of commands (or without built-in data structure locking), this property is crucial to preventing the inspection of intermediate (potentially inconsistent) states. Constant step (CSc) simulations are allowed a constant number of commands for each simulated command, and are thus appropriate when the state can grow to be large but the deployment scenario requires that the number of steps for any simulated action remain bounded (e.g., to prevent starvation due to locked structures).

**CT: Trace structure** properties restrict the path that the simulating system can take during the simulation of a single command. Semantic lock-step (CT1, depicted in Fig. 3) provides the benefits of a lock-step simulation in a slightly relaxed way: a “setup” phase prepares for the transition by changing only internal data (i.e., while remaining equivalent to the start state), then the transition occurs to a state equivalent to the end state, and then the “cleanup” phase cleans up any unnecessary leftover data (again, while remaining equivalent to the same end state). This is particularly useful when lock-step is too strict, but the deployment scenario is sensitive to the exposure of intermediate states (since, in CT1, no states are exposed except those equivalent to the start and end states). Query monotonicity (CTq) ensures that no query changes its truth value except those that are required to change between the start and end state. This allows multiple steps, but ensures that intermediate states, while not corresponding with the start or end state, never answer any query in a way that neither the start nor end state would. This is useful in scenarios where intermediate states are undesirable, but users are not expected to execute more than a single query between “valid” states (and will thus never detect the inconsistency). Access monotonicity (CTa) is similar, but applies only to authorization requests, and is useful in scenarios where inconsistent states are not a danger as long as they do not wrongly allow or forbid a request. Finally, non-contamination (CTs) ensures that no two accesses are allowed in an intermediate state that are not *both* allowed in either the start or end state. Thus, the simulating system is restricted not only from allowing accesses forbidden in the simulated system, but also combinations of individually-allowed accesses that are never combined in the simulated system. This restriction is particularly useful in environments with operations that “swap” accesses from one subject or object to another, or where separation of privilege is utilized.

**CA: Actor preservation** restricts which users can be invoked to simulate commands. Self-execution (CAT) requires each simulating command be executed by the same user as the original, simulated command. This allows the simulating agent to be completely unprivileged, mapping commands as a service to the user, but without executing them with any privilege beyond the user’s own. Administration-preservation (CAa) requires any non-administrative simulated command be mapped to a sequence of non-administrative commands (i.e., a command that does not invoke administrative privileges cannot be simulated by an administrative command). This corresponds to scenarios in which users will be expected to operate largely without administrative intervention. No restriction in this dimension means that the command mapping can return commands to be executed by any other user. This is most appropriate when the simulating agent is trusted to execute administrative actions on behalf of untrusted users, or when the commands returned can then be delegated to other users to be approved and executed.

Finally, **R: reachability** specifies whether the simulating system should be restricted from entering a state that does not correspond to a simulated state. If the simulation agent is




Fig. 6: Partial lattice of canonical usage

users’ only interface to the deployed access control system, forward reachability ( $R \rightarrow$ ) is sufficient. However, if users can access the simulating system’s native commands, bireachability ( $R \leftrightarrow$ ) ensures that the system cannot transition to a state that is inconsistent with the simulated system.

### C. Studying Canonical Usages

Next, we use the above interpretation of our expressiveness simulation properties to guide a discussion about how each of the notions of simulation that we studied in Section V is used by its creators. In many cases, the definition for a particular notion of simulation is underconstrained, and the simulations written within the framework actually satisfy stronger properties than the defined lower bound. We refer to the set of properties that the authors seem to intend for a simulation to uphold as its *canonical usage*. In the case of Sandhu’s simulation, the author recognizes that the given constructions are stronger than the definition, noting that formalizing the definition of the stronger simulation is beyond the scope of the work [2]. Here, we make conjectures regarding the decomposition of the canonical usage of these simulations. A lattice view of these conjectures is shown in Fig. 6, where  $\overline{\mathcal{X}}$  indicates the canonical usage of simulation type  $\mathcal{X}$ . For example,  $\overline{\text{SM}}$  refers to the form of the SMG simulation used in [7], [8].

It is interesting to note that the relationships between notions of simulation are not necessarily preserved in the canonical usage. While SMG by definition is the weakest simulation, the canonical usage  $\overline{\text{SM}}$  is incomparable to any simulation’s definition and positioned strictly weaker than the canonical usage of the CDM simulations. While, by definition, the TL state-matching reduction is more strict than HMG+ parameterized expressiveness, their canonical usages are incomparable due to  $\overline{\text{TL-SMR}}$  enforcing bireachability ( $R \leftrightarrow$ ) and using polynomial state size (SSp), compared to  $\overline{\text{HMG+}}$  enforcing forward reachability ( $R \rightarrow$ ) and using linear state size (SSl). Finally, we note that all of CDMs, CDMw, SMG, and ALS simulations are canonically used in such a way that enforces full query preservation (QPf); that is, all of the constructed mappings of these types use the identity mapping for all supported queries, despite the fact that none of them



specifically require this by definition. This trend of a notion of simulation's usage being consistently more strict than its definition reveals the difficulty in fully specifying the set of properties that a notion of expressiveness simulation is intended to enforce. The discussion in this section, aimed at helping analysts choose a reasonable set of properties for a deployment, can also help ensure that newer notions of simulation are fully specified, and best match their intended usages.

## VII. CONCLUSION AND FUTURE WORK

In this paper, we organize the existing knowledge of expressiveness simulations by formalizing a granular, property-based representation, proposing a wide range of dimensions of simulation properties, and positioning influential notions of expressiveness simulation from the literature within the lattice of these properties. In doing so, we provide the first systematic comparison of existing simulations that were not previously known to be directly comparable, showing how these notions of expressiveness simulation relate to one another.

Looking away from existing notions of simulation and rather *between* them, this work allows us to explore an organized space of simulations to identify areas to explore in future research. For instance, knowing expressiveness results derived using the SMG and ALS simulations, which of these hold true for notions of simulation “between” the two existing notions? What results can be shown for a simulation decomposing to the union of the properties of two existing notions? How far up the lattice do all systems become incomparable? These questions can only be explored thanks to the systematic means of simulation decomposition.

Finally, understanding the systems implications of various simulation properties will enable analysts to select the notion of access control expressiveness that corresponds most closely to the scenario in which they plan to deploy the target access control system(s). Thus, we make inroads toward bringing expressiveness analysis techniques out of the strictly formal realm, and repurpose these techniques to help select the most suitable access control system for a given application.

A question to be explored in future work is the identification of the set of analysis questions that a particular set of simulation properties preserve. For example, Tripunitara and Li showed that the state-matching reduction preserves *compositional security analysis instances*: the set of questions containing a single quantifier ( $\exists$  or  $\forall$ ), a propositional formula over queries  $\varphi$ , and a start state  $\gamma$  [13]. Semantically, the question asks whether  $\varphi$  it is {ever, always} true in states reachable from  $\gamma$ . If  $\mathcal{T}$  admits a state-matching reduction of  $\mathcal{S}$ , then all compositional security analysis instances have the same value in  $\mathcal{S}$  and  $\mathcal{T}$ . Identifying the types of analysis questions preserved by other notions of simulation would allow us even greater understanding of the practical and theoretical impacts of simulation property choices.

## REFERENCES

- [1] P. Ammann, R. J. Lipton, and R. S. Sandhu, “The expressive power of multi-parent creation in a monotonic access control model,” in *5th IEEE Computer Security Foundations Workshop - CSFW'92, Franconia, New Hampshire, USA, June 16-18, 1992, Proceedings*, Jun. 1992, pp. 148–156.
- [2] R. S. Sandhu, “Expressive power of the schematic protection model,” *Journal of Computer Security*, vol. 1, no. 1, pp. 59–98, 1992.
- [3] R. S. Sandhu and S. Ganta, “On testing for absence of rights in access control models,” in *CSFW*, 1993, pp. 109–118.
- [4] —, “On the expressive power of the unary transformation model,” in *Third European Symposium on Research in Computer Security*, Nov. 1994, pp. 301–318.
- [5] S. Ganta, “Expressive power of access control models based on propagation of rights,” Ph.D. dissertation, George Mason University, 1996.
- [6] P. Ammann, R. J. Lipton, and R. S. Sandhu, “The expressive power of multi-parent creation in monotonic access control models,” *Journal of Computer Security*, vol. 4, no. 2/3, pp. 149–166, 1996.
- [7] R. S. Sandhu and Q. Munawar, “How to do discretionary access control using roles,” in *ACM Workshop on Role-Based Access Control*, 1998, pp. 47–54.
- [8] Q. Munawar and R. S. Sandhu, “Simulation of the augmented typed access matrix model (ATAM) using roles,” in *INFOSEC99*, 1999.
- [9] S. L. Osborn, R. S. Sandhu, and Q. Munawar, “Configuring role-based access control to enforce mandatory and discretionary access control policies,” *ACM Trans. Inf. Syst. Secur.*, vol. 3, no. 2, pp. 85–106, 2000.
- [10] A. Chander, D. Dean, and J. C. Mitchell, “A state-transition model of trust management and access control,” in *CSFW*, 2001, pp. 27–43.
- [11] E. Bertino, B. Catania, E. Ferrari, and P. Perlasca, “A logical framework for reasoning about access control models,” *ACM Trans. Inf. Syst. Secur.*, vol. 6, no. 1, pp. 71–127, 2003.
- [12] M. V. Tripunitara and N. Li, “Comparing the expressive power of access control models,” in *Proceedings of the 11th ACM Conference on Computer and Communications Security, CCS 2004, Washington, DC, USA, October 25-29, 2004*, Oct. 2004, pp. 62–71.
- [13] —, “A theory for comparing the expressive power of access control models,” *Journal of Computer Security*, vol. 15, no. 2, pp. 231–272, 2007.
- [14] T. L. Hinrichs, D. Martinoia, W. C. Garrison III, A. J. Lee, A. Panebianco, and L. Zuck, “Application-sensitive access control evaluation using parameterized expressiveness,” in *Proceedings of the 26th IEEE Computer Security Foundations Symposium (CSF)*, Jun. 2013, pp. 145–160.
- [15] W. C. Garrison III, Y. Qiao, and A. J. Lee, “On the suitability of dissemination-centric access control models for group-centric sharing,” in *Proceedings of the Fourth ACM Conference on Data and Application Security and Privacy*, Mar. 2014, pp. 1–12.
- [16] W. C. Garrison III, A. J. Lee, and T. L. Hinrichs, “An actor-based, application-aware access control evaluation framework,” in *19th ACM Symposium on Access Control Models and Technologies*, Jun. 2014, pp. 199–210.
- [17] M. A. Harrison, W. L. Ruzzo, and J. D. Ullman, “Protection in operating systems,” *Communications of the ACM*, vol. 19, no. 8, pp. 461–471, Aug. 1976.
- [18] N. Li and M. V. Tripunitara, “Security analysis in role-based access control,” *ACM Trans. Inf. Syst. Secur.*, vol. 9, no. 4, pp. 391–420, 2006.
- [19] T. L. Hinrichs, W. C. Garrison III, A. J. Lee, S. Saunders, and J. C. Mitchell, “TBA : A hybrid of logic and extensional access control systems,” in *Formal Aspects of Security and Trust, 8th International Workshop*, Sep. 2011, pp. 198–213.
- [20] P. Naldurg and R. H. Campbell, “Dynamic access control: preserving safety and trust for network defense operations,” in *SACMAT 2003, 8th ACM Symposium on Access Control Models and Technologies, June 2-3, 2003, Villa Gallia, Como, Italy, Proceedings*, 2003, pp. 231–237.
- [21] T. Bourdier, H. Cirstea, M. Jaume, and H. Kirchner, “Formal specification and validation of security policies,” in *Foundations and Practice of Security - 4th Canada-France MITACS Workshop, FPS 2011, Paris, France, May 12-13, 2011, Revised Selected Papers*, 2011, pp. 148–163.
- [22] B. D. Payne, R. Sailer, R. Cáceres, R. Perez, and W. Lee, “A layered approach to simplified access control in virtualized systems,” *Operating Systems Review*, vol. 41, no. 4, pp. 12–19, 2007.
- [23] R. Krishnan, R. S. Sandhu, J. Niu, and W. H. Winsborough, “Foundations for group-centric secure information sharing models,” in *SACMAT 2009, 14th ACM Symposium on Access Control Models and Technologies, Stresa, Italy, June 3-5, 2009, Proceedings*, 2009, pp. 115–124.
- [24] V. C. Hu, D. A. Frincke, and D. F. Ferraio, “The policy machine for security policy management,” in *Computational Science - ICCS 2001*,

International Conference, San Francisco, CA, USA, May 28-30, 2001. *Proceedings, Part II*, May 2001, pp. 494–506.

- [25] D. F. Ferraiolo, S. I. Gavrila, V. C. Hu, and D. R. Kuhn, “Composing and combining policies under the policy machine,” in *SACMAT 2005, 10th ACM Symposium on Access Control Models and Technologies, Stockholm, Sweden, June 1-3, 2005, Proceedings*, Jun. 2005, pp. 11–20.
- [26] V. C. Hu, D. F. Ferraiolo, and D. R. Kuhn, *Assessment of Access Control Systems*. National Institute of Standards and Technology, 2006.
- [27] K. Kane and J. C. Browne, “On classifying access control implementations for distributed systems,” in *SACMAT 2006, 11th ACM Symposium on Access Control Models and Technologies, Lake Tahoe, California, USA, June 7-9, 2006, Proceedings*, Jun. 2006, pp. 29–38.
- [28] O. Goldreich, *Computational Complexity: A Conceptual Perspective*. Cambridge University Press, 2008.

## APPENDIX A

### QUERY DECIDER PROPERTIES

In this appendix, we present the dimensions in which the query decider can be restricted. Recall (Definition 5) that the query decider is a function  $\sigma_Q = Q^S \times \Gamma^T \rightarrow \{\text{TRUE}, \text{FALSE}\}$  that assigns a truth value to each query in system  $\mathcal{S}$ , in each state in system  $\mathcal{T}$ . As with the command mapping, we can restrict what properties of the state the query decider can observe, but in general it sees the whole state. Thus, it allows us to answer  $\mathcal{S}$  queries in an active simulation using  $\mathcal{T}$ . We now discuss the ways in which we can restrict this decider.

#### Dimension QD: Query decider dependence

While Definition 5 maps each  $\mathcal{S}$  query and  $\mathcal{T}$  state to a truth value, previous works use more strict query deciders, ranging from mapping each  $\mathcal{S}$  query  $q$  to a  $\mathcal{T}$  query  $q'$  and returning  $\gamma^T \vdash q'$  [13], to mapping each  $\mathcal{S}$  query and  $\mathcal{T}$  theory to a truth value [14] (calculating the truth value by observing only the queriable portions of the  $\mathcal{T}$  state). *Query decider dependence*, similar to command mapping dependence, restricts the information that the query decider can consider about a  $\mathcal{T}$  state when deciding the truth value of an  $\mathcal{S}$  query in that state.

##### QD1: Independent, unitary-range query decider

$$\exists \sigma' : Q^S \rightarrow Q^T. (\sigma_Q(q, \gamma) \equiv \gamma \vdash \sigma'(q))$$

##### QDi: Independent query decider

$$\exists \sigma' : Q^S \rightarrow \mathcal{L}(Q^T). (\sigma_Q(q, \gamma) \equiv \gamma \vdash \sigma'(q))$$

##### QDt: Theory-dependent query decider

$$\exists \sigma' : Q^S \times Th(\mathcal{T}) \rightarrow \{\text{TRUE}, \text{FALSE}\}. \sigma_Q(q, \gamma) \equiv \sigma'(q, Th(\gamma))$$

##### QDs: State-dependent query decider *No restriction*.

First, *independent, unitary-range query* simulations first map each  $\mathcal{S}$  query  $q^S$  to a single  $\mathcal{T}$  query  $q^T$ , then return the truth value of asking this  $\mathcal{T}$  query in  $\gamma^T$  (i.e., returns  $\gamma^T \vdash q^T$ ). This type of simulation was used in [13]. *Independent query* simulations, similarly, map  $q^S$  independent of the state, but in this case can map it to an element of  $\mathcal{L}(Q^T)$ , a boolean expression over  $Q^T$ , rather than a single query. This allows the simulation, for example, to map the  $\mathcal{S}$  query “Is user  $u$  a member of role  $r$ ” to  $\mathcal{T}$  query sentence “Does user  $u$  exist and does user  $u$  have attribute role: $r$ ?”

*Theory-dependent query* allows the truth value to be determined based only on the *theory* of the  $\mathcal{T}$  state, or the values

of all  $\mathcal{T}$  queries in the state. This allows the query decider to inspect the values of (potentially) all  $\mathcal{T}$  queries, but does not allow the decider to consider any features of the state that cannot be queried using  $Q^T$ . This version of the query decider is used in [14]. Finally, we refer to the general case as *state-dependent query*, since in this case the query decider can arbitrarily inspect the  $\mathcal{T}$  state before returning a truth value for an  $\mathcal{S}$  query, rather than being restricted only to those elements of the state which are observable by asking queries in  $Q^T$ .

The biggest impact a selection in QD has is to limit the privilege of the simulating agent: under QD1 and QDi, the simulating agent need know nothing about the current state of the system to map queries, and simulated queries are mapped statically to dynamic queries. Under QDt, the simulating agent must be privileged to view the values of all queries at runtime, while QDs assumes the ability to arbitrarily inspect the state at runtime (even the portions of state that are not queriable).

#### Dimension QC: Query decider complexity

As with the command mapping, due to potential resource constraints, we may enforce limits on the runtime complexity of the query decider routine, with respect to the size of the state it is executed in. As with commands, this is not applicable for independent query deciders.

**QCc: Constant query decider**  $\forall q \in Q^S$ , the algorithm for  $\sigma_q(\gamma) = \sigma_Q(q, \gamma)$  has time complexity  $T(n) \in \mathcal{O}(1)$

**QC $\infty$ : Unbounded query decider** *No restriction*.

For theory-dependent and state-dependent query deciders, we can limit the complexity of the procedure—here, we consider only the constant-time restriction explicitly, though other restrictions may be useful in some cases.

#### Dimension QP: Query preservation

Query preservation is a property dimension that indicates which queries need to stay the same as they are mapped from system  $\mathcal{S}$  to system  $\mathcal{T}$ . A particular application may require any given set of queries to be preserved; here, we present several generic examples.

##### QPF: Complete query preservation

$$\forall q \in Q^S, \gamma \in \Gamma^T : \sigma_Q(q, \gamma) \equiv \gamma \vdash q$$

##### QPa: Authorization preservation

$$\forall r \in \mathcal{R}^S, \gamma \in \Gamma^T : \sigma_Q(r, \gamma) \equiv \gamma \vdash r$$

##### QPw: Weak authorization preservation

$\exists f : \mathcal{R}^S \rightarrow \mathcal{R}^T$  such that the following two conditions hold:

- $\forall r \in \mathcal{R}^S, \gamma \in \Gamma^T. (\sigma_Q(r, \gamma) = \text{TRUE} \Rightarrow \gamma \vdash f(r))$
- $\forall r \in \mathcal{R}^T, \gamma \in \Gamma^T. (\gamma \vdash r \Rightarrow \exists r^S. (\sigma_Q(r^S, \gamma) = \text{TRUE} \wedge f(r^S) = r))$

The most common property in this dimension is *authorization preservation*, which roughly enforces that the query decider maps each  $\mathcal{S}$  request to the value of the identical request in the  $\mathcal{T}$  state. This requires  $\mathcal{T}$  to accept at least the same authorization requests as  $\mathcal{S}$ , and can be seen as ensuring



that  $\mathcal{T}$  is using its model “as intended” (i.e., forcing it to answer simulated requests as it would its own native requests). Complete query preservation restricts the query decider in the same way, but for all supported queries.

Authorization preservation was formalized in [14] (as AC-preservation), but has been used implicitly in other simulations (e.g., [10]) that do not include any mapping from  $\mathcal{S}$  requests to  $\mathcal{T}$  requests (i.e., assume the identity mapping). For formalisms that do not consider queries other than requests, authorization preservation and complete query preservation are equivalent.

A related property is *weak authorization preservation*, defined in [15] (as weak AC-preservation). This property is a weakened version of authorization preservation: its intentions are similar, but the weak form can be used even when  $\mathcal{S}$  and  $\mathcal{T}$  do not support the exact same requests (i.e., simulating a system with requests of the form, “Does user  $u$  have access to permission  $p$ ?” in a system with requests of the form, “Does subject  $s$  have access *read* to object  $o$ ?”). Weak authorization preservation allows a request transformation function, which maps  $\mathcal{S}$  requests to  $\mathcal{T}$  requests. The definition of this property ensures that each  $\mathcal{S}$  request is mapped into  $\mathcal{T}$ , and each  $\mathcal{T}$  request that is granted represents some  $\mathcal{S}$  request.

#### A. Interactions Between Dimensions

Independence query decider (QDi) implies constant query decider (QCc), since a query decider that does not depend on the state must be constant-time with respect to the state size. Full query preservation (QPf) implies constant query decider (QCc), since the identity mapping is a constant-time procedure, and implies unitary-range query decider (QD1), since the identity mapping always outputs only one query.

#### B. Interpreting the Dimensions

**QD: Query dependence**, similar to command dependence, can restrict the query decider from using the full knowledge of the state in question when mapping queries to truth values. Independent query decider (QD1 and QDi) map each simulated query to its simulating queries, and its truth value is then determined by checking the values of these mapped queries in the simulating state. This is useful, as with commands, for precompilation, and restricts the simulating agent’s role to be akin to a Karp reduction (i.e., it can only return the value of its simulating queries with no modifications) [28, p. 60]. Theory-dependent query decider (QDt) allows the inspection of the full theory of the state, and thus arbitrary computation over the values of all queries in the simulating system, corresponding to a query decider that is only privileged enough to ask queries, but not inspect state. Finally, state-dependent query decider (QDs) allows arbitrary inspection of the state, and thus may answer simulated queries using state that is usually unqueriable in the simulating system; this requires a powerful simulation agent that is trusted to view the full access control data structures.

**QC: Query complexity** restricts the time-complexity of the query decider with respect to the size of the simulating state, and is thus chosen for reasons analogous to command complexity (CC).

**QP: Query preservation** restricts certain queries to be mapped by a sort of “identity function.” That is, certain simulated queries are mapped TRUE if and only if the query is also present in the simulating system, and the query returns TRUE in the current simulating state. This dimension of restrictions is strict in that it requires the simulating system to support all of the included queries of the simulated system, but ensures that the simulating system is used as per normal. If the deployment scenario is unable to cope with reduced throughput caused by the query decider during simulation, QPf also ensures that we can pipe simulated queries directly to the simulating system. Particularly useful is authorization preservation (QP<sub>a</sub>) when using a simulating system based on a model with formally-proven properties; Since all authorization requests must be mapped by the identity, the simulating system allows a simulated request exactly when it would natively allow the same request.

### APPENDIX B DECOMPOSITION PROOFS

#### A. TL State-Matching Reduction

For the purposes of this proof, let the set of simulation properties  $\mathcal{P} = \{\text{SCq}, \text{QD1}, \text{R}\leftrightarrow\}$ .

**Lemma 4** *Given access control systems  $\mathcal{S}$ ,  $\mathcal{T}$ , and  $\mathcal{U}$ ,*

$$\mathcal{T} \xrightarrow[\text{TL-SMR}]{\text{sim}} \mathcal{S} \wedge \mathcal{S} \xrightarrow[\mathcal{P}]{\text{sim}} \mathcal{U} \Rightarrow \mathcal{T} \xrightarrow[\mathcal{P}]{\text{sim}} \mathcal{U}$$

*That is, if  $\mathcal{T}$  admits a state-matching reduction of  $\mathcal{S}$ , and  $\mathcal{S}$  admits a simulation of  $\mathcal{U}$  with properties  $\{\text{SCq}, \text{QD1}, \text{R}\leftrightarrow\}$ , then  $\mathcal{T}$  admits a simulation of  $\mathcal{U}$  with properties  $\{\text{SCq}, \text{QD1}, \text{R}\leftrightarrow\}$ .*

*Proof:* Let  $\mathcal{S}$ ,  $\mathcal{T}$ , and  $\mathcal{U}$  be arbitrary access control systems such that  $\mathcal{T} \xrightarrow[\text{TL-SMR}]{\text{sim}} \mathcal{S}$  and  $\mathcal{S} \xrightarrow[\mathcal{P}]{\text{sim}} \mathcal{U}$ . To prove Lemma 4, we must then show that  $\mathcal{T} \xrightarrow[\mathcal{P}]{\text{sim}} \mathcal{U}$ .

Since  $\mathcal{S} \xrightarrow[\mathcal{P}]{\text{sim}} \mathcal{U}$ , by QD1,

$$\exists \sigma^{\text{QD1}} : \mathcal{Q}^{\mathcal{U}} \rightarrow \mathcal{Q}^{\mathcal{S}}. (\sigma_Q(q^{\mathcal{U}}, \gamma^{\mathcal{S}}) \equiv \gamma^{\mathcal{S}} \vdash \sigma^{\text{QD1}}(q^{\mathcal{U}}))$$

Since  $\mathcal{T} \xrightarrow[\text{TL-SMR}]{\text{sim}} \mathcal{S}$ , the state-matching reduction provides a mapping from  $\mathcal{Q}^{\mathcal{S}}$  to  $\mathcal{Q}^{\mathcal{T}}$ . Call this mapping  $\sigma^{\text{SMR}}$ .

Thus, let  $\sigma' : \mathcal{Q}^{\mathcal{U}} \rightarrow \mathcal{Q}^{\mathcal{T}} = \sigma^{\text{SMR}} \cdot \sigma^{\text{QD1}}$ , and say  $\sigma_Q(q^{\mathcal{U}}, \gamma^{\mathcal{T}}) \equiv \gamma^{\mathcal{T}} \vdash \sigma'(q^{\mathcal{U}})$ . This forms a query decider that satisfies QD1.

Choose an arbitrary state  $\gamma_0^{\mathcal{U}} \in \Gamma^{\mathcal{U}}$  and command  $\psi^{\mathcal{U}} \in \Psi^{\mathcal{U}}$ , and let  $\text{next}(\gamma_0^{\mathcal{U}}, \psi^{\mathcal{U}}) = \gamma_1^{\mathcal{U}}$ . Let  $\gamma_0^{\mathcal{S}} \in \Gamma^{\mathcal{S}}$  such that  $\gamma_0^{\mathcal{U}} \stackrel{\mathcal{Q}}{\sim} \gamma_0^{\mathcal{S}}$ . Since  $\mathcal{S} \xrightarrow[\mathcal{P}]{\text{sim}} \mathcal{U}$ ,

$$\exists \gamma_1^{\mathcal{S}} \in \Gamma^{\mathcal{S}}. (\text{terminal}(\gamma_0^{\mathcal{S}}, \sigma_{\Psi}(\psi^{\mathcal{U}}, \gamma_0^{\mathcal{S}})) = \gamma_1^{\mathcal{S}} \wedge \gamma_1^{\mathcal{U}} \stackrel{\mathcal{Q}}{\sim} \gamma_1^{\mathcal{S}})$$

Let  $\gamma_0^{\mathcal{T}} \in \Gamma^{\mathcal{T}}$  such that  $\gamma_0^{\mathcal{S}} \stackrel{\mathcal{Q}}{\sim} \gamma_0^{\mathcal{T}}$ . Since  $\mathcal{T} \xrightarrow[\text{TL-SMR}]{\text{sim}} \mathcal{S}$ ,

$$\exists \gamma_1^{\mathcal{T}} \in \Gamma^{\mathcal{T}}. (\gamma_0^{\mathcal{T}} \xrightarrow{*} \gamma_1^{\mathcal{T}} \wedge \gamma_1^{\mathcal{S}} \stackrel{\mathcal{Q}}{\sim} \gamma_1^{\mathcal{T}})$$

Thus, there exists a sequence of  $\mathcal{T}$  commands  $\Psi_0^{\mathcal{T}}$  such that  $\text{terminal}(\gamma_0^{\mathcal{T}}, \Psi_0^{\mathcal{T}}) = \gamma_1^{\mathcal{T}}$ . Define  $\sigma_{\Psi} : \Psi^{\mathcal{U}} \times \Gamma^{\mathcal{T}} \rightarrow (\Psi^{\mathcal{T}})^*$  such that it returns  $\Psi_0^{\mathcal{T}}$  for  $\gamma_0^{\mathcal{T}}, \psi^{\mathcal{U}}$ . This is formed by concatenating a sequence of sequences of commands: for

each command  $\psi_i^S$  that  $\mathcal{S}$  needs to execute to simulate  $\psi^U$ , concatenate the commands that  $\mathcal{T}$  needs to execute to simulate  $\psi_i^S$ .

Then, given  $\gamma_0^U, \gamma_1^U \in \Gamma^U, \gamma_0^T \in \Gamma^T, \psi^U \in \Psi^U$  such that  $\text{next}(\gamma_0^U, \psi^U) = \gamma_1^U$ , and  $\gamma_0^U \stackrel{q}{\sim} \gamma_0^T$ ,

$$\exists \gamma_1^T \in \Gamma^T. (\text{terminal}(\gamma_0^T, \sigma_\Psi(\psi, \gamma_0^T)) = \gamma_1^T \wedge \gamma_1^S \stackrel{q}{\sim} \gamma_1^T)$$

Hence,  $\mathcal{T} \stackrel{\text{sim}}{\sim}_{\{\text{SCq}, \text{QD1}, \text{R} \leftrightarrow\}} \mathcal{U}$ . Next we show  $\text{R} \leftrightarrow$ .

Choose some arbitrary states  $\gamma_0^T, \gamma_1^T \in \Gamma^T$  such that  $\gamma_0^T \mapsto \gamma_1^T$ . Let  $\gamma_0^S \in \Gamma^S$  such that  $\gamma_0^S \stackrel{q}{\sim} \gamma_0^T$ . Since  $\mathcal{T} \stackrel{\text{sim}}{\sim}_{\text{TL-SMR}} \mathcal{S}$ ,

$$\exists \gamma_1^S. (\gamma_0^S \mapsto \gamma_1^S \wedge \gamma_1^S \stackrel{q}{\sim} \gamma_1^T)$$

Let  $\gamma_0^U \in \Gamma^U$  such that  $\gamma_0^U \stackrel{q}{\sim} \gamma_0^S$ . Since  $\mathcal{S} \stackrel{\text{sim}}{\sim}_{\mathcal{P}} \mathcal{U}$ ,

$$\exists \gamma_1^U. (\gamma_0^U \mapsto \gamma_1^U \wedge \gamma_1^U \stackrel{q}{\sim} \gamma_1^S)$$

Thus, given  $\gamma_0^T, \gamma_1^T \in \Gamma^T, \gamma_0^U \in \Gamma^U$  such that  $\gamma_0^T \mapsto \gamma_1^T$  and  $\gamma_0^U \stackrel{q}{\sim} \gamma_0^T$ ,

$$\exists \gamma_1^U \in \Gamma^U. (\gamma_0^U \mapsto \gamma_1^U \wedge \gamma_1^U \stackrel{q}{\sim} \gamma_1^T)$$

Hence,  $\mathcal{T} \stackrel{\text{sim}}{\sim}_{\mathcal{P}} \mathcal{U}$ .  $\blacksquare$

**Lemma 5** Given access control systems  $\mathcal{S}$  and  $\mathcal{T}$  and simulation properties  $\mathcal{P} = \{\text{SCq}, \text{QD1}, \text{R} \leftrightarrow\}$ ,  $\mathcal{S} \leq^{\mathcal{P}} \mathcal{T} \Rightarrow \mathcal{T} \stackrel{\text{sim}}{\sim}_{\text{TL-SMR}} \mathcal{S}$ . That is, if  $\mathcal{T}$  is at least as expressive as  $\mathcal{S}$  with respect to properties  $\mathcal{P}$ , then  $\mathcal{T}$  admits a state-matching reduction of  $\mathcal{S}$ .

*Proof:* Let  $\mathcal{S}$  and  $\mathcal{T}$  be arbitrary access control systems such that  $\mathcal{S} \leq^{\mathcal{P}} \mathcal{T}$ . Since  $\mathcal{S} \leq^{\mathcal{P}} \mathcal{T}$ , for any access control system  $\mathcal{U}$ , if  $\mathcal{S} \stackrel{\text{sim}}{\sim}_{\mathcal{P}} \mathcal{U}$ , then  $\mathcal{T} \stackrel{\text{sim}}{\sim}_{\mathcal{P}} \mathcal{U}$ .

Since  $\mathcal{S}$  can trivially simulate itself,  $\mathcal{S} \stackrel{\text{sim}}{\sim}_{\mathcal{P}} \mathcal{S}$ , and thus  $\mathcal{T} \stackrel{\text{sim}}{\sim}_{\mathcal{P}} \mathcal{S}$ .

By QD1,  $\exists \sigma' : Q^S \rightarrow Q^T. (\sigma_Q(q^S, \gamma^T) \equiv \gamma^T \vdash \sigma'(q^S))$ . Thus,  $\sigma'$  satisfies the format of the TL-SMR query mapping (i.e.,  $\sigma_Q : Q^S \rightarrow Q^T$ ).

Then, given  $\gamma_0^S, \gamma_1^S \in \Gamma^S, \gamma_0^T \in \Gamma^T$ , by  $\text{R} \rightarrow$ , if  $\gamma_0^S \stackrel{q}{\sim} \gamma_0^T$  and  $\gamma_0^S \mapsto \gamma_1^S$ , then

$$\exists \gamma_1^T. (\gamma_0^T \mapsto \gamma_1^T \wedge \gamma_1^S \stackrel{q}{\sim} \gamma_1^T)$$

Since SCq satisfies the TL-SMR definition of state correspondence, this means we have satisfied the first property of the state-matching reduction.

1) For every state  $\gamma_1^S$  in system  $\mathcal{S}$  such that  $\gamma_0^S \mapsto \gamma_1^S$ , there exists a state  $\gamma_1^T$  such that  $\gamma_0^T \mapsto \gamma_1^T$  and  $\gamma_1^S$  and  $\gamma_1^T$  are equivalent under  $\sigma$ .

And by bidirectional reachability, given  $\gamma_0^S \in \Gamma^S, \gamma_0^T, \gamma_1^T \in \Gamma^T$ , if  $\gamma_0^S \stackrel{q}{\sim} \gamma_0^T$  and  $\gamma_0^T \mapsto \gamma_1^T$ , then

$$\exists \gamma_1^S. (\gamma_0^S \mapsto \gamma_1^S \wedge \gamma_1^S \stackrel{q}{\sim} \gamma_1^T)$$

And therefore, we have satisfied the second property of the state-matching reduction:

2) For every state  $\gamma_1^T$  in system  $\mathcal{T}$  such that  $\gamma_0^T \mapsto \gamma_1^T$ , there exists a state  $\gamma_1^S$  such that  $\gamma_0^S \mapsto \gamma_1^S$  and  $\gamma_1^T$  and  $\gamma_1^S$  are equivalent under  $\sigma$ .

These properties satisfy the definition for a state-matching reduction, and hence  $\mathcal{T}$  admits a state-matching reduction of  $\mathcal{S}$  ( $\mathcal{T} \stackrel{\text{sim}}{\sim}_{\text{TL-SMR}} \mathcal{S}$ ).  $\blacksquare$

**Theorem 6**  $\text{TL-SMR} \doteq \{\text{SCq}, \text{QD1}, \text{R} \leftrightarrow\}$ ; that is, the TL state-matching reduction decomposes to query correspondence; independent, unitary-range query; and bidirectional reachability.

*Proof:* By Lemma 4, if  $\mathcal{T} \stackrel{\text{sim}}{\sim}_{\text{TL-SMR}} \mathcal{S}$ , then  $\mathcal{S} \leq^{\mathcal{P}} \mathcal{T}$ . By Lemma 5, if  $\mathcal{S} \leq^{\mathcal{P}} \mathcal{T}$ , then  $\mathcal{T} \stackrel{\text{sim}}{\sim}_{\text{TL-SMR}} \mathcal{S}$ . Thus,  $\mathcal{S} \leq^{\mathcal{P}} \mathcal{T}$  if and only if  $\mathcal{T} \stackrel{\text{sim}}{\sim}_{\text{TL-SMR}} \mathcal{S}$ , and thus the state-matching reduction decomposes to  $\{\text{SCq}, \text{QD1}, \text{R} \leftrightarrow\}$ .  $\blacksquare$

### B. HMG+ Parameterized Expressiveness Simulation

For the purposes of this proof, let the set of simulation properties  $\mathcal{P} = \{\text{SCq}, \text{QDt}, \text{R} \leftrightarrow\}$ .

**Lemma 7** Given access control systems  $\mathcal{S}$ ,  $\mathcal{T}$ , and  $\mathcal{U}$ ,

$$\mathcal{T} \stackrel{\text{sim}}{\sim}_{\text{HMG}+} \mathcal{S} \wedge \mathcal{S} \stackrel{\text{sim}}{\sim}_{\mathcal{P}} \mathcal{U} \Rightarrow \mathcal{T} \stackrel{\text{sim}}{\sim}_{\mathcal{P}} \mathcal{U}$$

That is, if  $\mathcal{T}$  admits an HMG+ simulation of  $\mathcal{S}$ , and  $\mathcal{S}$  admits a simulation of  $\mathcal{U}$  with properties  $\{\text{SCq}, \text{QDt}, \text{R} \leftrightarrow\}$ , then  $\mathcal{T}$  admits a simulation of  $\mathcal{U}$  with properties  $\{\text{SCq}, \text{QDt}, \text{R} \leftrightarrow\}$ .

*Proof:* Let  $\mathcal{S}$ ,  $\mathcal{T}$ , and  $\mathcal{U}$  be arbitrary access control systems such that  $\mathcal{T} \stackrel{\text{sim}}{\sim}_{\text{HMG}+} \mathcal{S}$  and  $\mathcal{S} \stackrel{\text{sim}}{\sim}_{\mathcal{P}} \mathcal{U}$ . To prove Lemma 7, we must then show that  $\mathcal{T} \stackrel{\text{sim}}{\sim}_{\mathcal{P}} \mathcal{U}$ .

Since  $\mathcal{S} \stackrel{\text{sim}}{\sim}_{\mathcal{P}} \mathcal{U}$ , by QDt,

$$\exists \sigma' : Q^U \times \text{Th}(\mathcal{S}) \rightarrow \{\text{TRUE}, \text{FALSE}\}. \sigma_Q(q, \gamma) \equiv \sigma'(q, \text{Th}(\gamma))$$

Since  $\mathcal{T} \stackrel{\text{sim}}{\sim}_{\text{HMG}+} \mathcal{S}$ , the HMG+ simulation contains the mapping  $\pi : Q^S \times \text{Th}(\mathcal{T}) \rightarrow \{\text{TRUE}, \text{FALSE}\}$ .

Thus, let  $\sigma'' : Q^U \times \text{Th}(\mathcal{T}) \rightarrow \{\text{TRUE}, \text{FALSE}\}$  be constructed as follows. Use  $\sigma' : Q^U \times \text{Th}(\mathcal{S})$ , and for each query  $q$  in  $\text{Th}(\mathcal{S})$  that is needed by  $\sigma'$ , consult  $\pi : Q^S \times \text{Th}(\mathcal{T})$  to obtain a truth value in the current  $\mathcal{T}$  state. This forms a query decider that satisfies QDt.

Choose an arbitrary state  $\gamma_0^U \in \Gamma^U$  and command  $\psi^U \in \Psi^U$ , and let  $\text{next}(\gamma_0^U, \psi^U) = \gamma_1^U$ . Let  $\gamma_0^S \in \Gamma^S$  such that  $\gamma_0^U \stackrel{q}{\sim} \gamma_0^S$ . Since  $\mathcal{S} \stackrel{\text{sim}}{\sim}_{\mathcal{P}} \mathcal{U}$ ,

$$\exists \gamma_1^S \in \Gamma^S. (\text{terminal}(\gamma_0^S, \sigma_\Psi(\psi^U, \gamma_0^S)) = \gamma_1^S \wedge \gamma_1^U \stackrel{q}{\sim} \gamma_1^S)$$

Let  $\gamma_0^T \in \Gamma^T$  such that  $\gamma_0^S \stackrel{q}{\sim} \gamma_0^T$ . Since  $\mathcal{T} \stackrel{\text{sim}}{\sim}_{\text{HMG}+} \mathcal{S}$ , we can map each command of the sequence  $\sigma_\Psi(\psi^U, \gamma_0^S)$  to a sequence of  $\mathcal{T}$  commands using the HMG+ simulation. Concatenating this sequence of sequences to a single sequence  $\Psi_0^T$ , and using HMG+ correctness:

$$\exists \gamma_1^T \in \Gamma^T. (\text{terminal}(\gamma_0^T, \Psi_0^T) = \gamma_1^T \wedge \gamma_1^S \stackrel{q}{\sim} \gamma_1^T)$$

Define  $\sigma_\Psi : \Psi^\mathcal{U} \times \Gamma^\mathcal{T} \rightarrow (\Psi^\mathcal{T})^*$  such that it returns  $\Psi_0^\mathcal{T}$  for  $\gamma_0^\mathcal{T}, \psi^\mathcal{U}$ .

Then, given  $\gamma_0^\mathcal{U}, \gamma_1^\mathcal{U} \in \Gamma^\mathcal{U}, \gamma_0^\mathcal{T} \in \Gamma^\mathcal{T}, \psi^\mathcal{U} \in \Psi^\mathcal{U}$  such that  $\text{next}(\gamma_0^\mathcal{U}, \psi^\mathcal{U}) = \gamma_1^\mathcal{U}$ , and  $\gamma_0^\mathcal{U} \stackrel{q}{\sim} \gamma_0^\mathcal{T}$ ,

$$\exists \gamma_1^\mathcal{T} \in \Gamma^\mathcal{T}. (\text{terminal}(\gamma_0^\mathcal{T}, \sigma_\Psi(\psi, \gamma_0^\mathcal{T})) = \gamma_1^\mathcal{T} \wedge \gamma_1^\mathcal{S} \stackrel{q}{\sim} \gamma_1^\mathcal{T})$$

Hence,  $\mathcal{T} \stackrel{\text{sim}}{\underset{\{\text{SCq}, \text{QDt}, \text{R} \rightarrow\}}{\sim}} \mathcal{U}$ . ■

**Lemma 8** Given access control systems  $\mathcal{S}$  and  $\mathcal{T}$  and simulation properties  $\mathcal{P} = \{\text{SCq}, \text{QDt}, \text{R} \rightarrow\}$ ,  $\mathcal{S} \leq^\mathcal{P} \mathcal{T} \Rightarrow \mathcal{T} \stackrel{\text{sim}}{\underset{\text{HMG}+}{\sim}} \mathcal{S}$ . That is, if  $\mathcal{T}$  is at least as expressive as  $\mathcal{S}$  with respect to properties  $\mathcal{P}$ , then  $\mathcal{T}$  admits an HMG+ simulation of  $\mathcal{S}$ .

*Proof:* Let  $\mathcal{S}$  and  $\mathcal{T}$  be arbitrary access control systems such that  $\mathcal{S} \leq^\mathcal{P} \mathcal{T}$ . Since  $\mathcal{S} \leq^\mathcal{P} \mathcal{T}$ , for any access control system  $\mathcal{U}$ , if  $\mathcal{S} \stackrel{\text{sim}}{\underset{\mathcal{P}}{\sim}} \mathcal{U}$ , then  $\mathcal{T} \stackrel{\text{sim}}{\underset{\mathcal{P}}{\sim}} \mathcal{U}$ .

Since  $\mathcal{S}$  can trivially simulate itself,  $\mathcal{S} \stackrel{\text{sim}}{\underset{\mathcal{P}}{\sim}} \mathcal{S}$ , and thus  $\mathcal{T} \stackrel{\text{sim}}{\underset{\mathcal{P}}{\sim}} \mathcal{S}$ .

By QDt,

$$\exists \sigma' : Q^\mathcal{S} \times \text{Th}(\mathcal{T}) \rightarrow \{\text{TRUE}, \text{FALSE}\}. \sigma_Q(q^\mathcal{S}, \gamma^\mathcal{T}) \equiv \sigma'(q^\mathcal{S}, \text{Th}(\gamma^\mathcal{T})) \Rightarrow \mathcal{T} \stackrel{\text{sim}}{\underset{\text{HMG}+}{\sim}} \mathcal{S}.$$

Thus,  $\sigma'$  satisfies the format of the HMG+ query mapping (i.e.,  $\pi : Q^\mathcal{S} \times \text{Th}(\mathcal{T}) \rightarrow \{\text{TRUE}, \text{FALSE}\}$ ), and by SCq, the state mapping preserves the query mapping, property (i) for the HMG+ correct simulation.

Let  $\gamma_0^\mathcal{S} \in \Gamma^\mathcal{S}, \psi \in \Psi^\mathcal{S}$  be an arbitrary state and command in  $\mathcal{S}$ , and  $\gamma_0^\mathcal{T} \in \Gamma^\mathcal{T}$  a state in  $\mathcal{T}$  such that  $\sigma_\Gamma(\gamma_0^\mathcal{S}) = \gamma_0^\mathcal{T}$ . Then, if  $\text{next}(\gamma_0^\mathcal{S}, \psi^\mathcal{S}) = \gamma_1^\mathcal{S}$ ,

$$\exists \gamma_1^\mathcal{T}. (\text{terminal}(\gamma_0^\mathcal{T}, \sigma_\Psi(\psi, \gamma_0^\mathcal{T})) = \gamma_1^\mathcal{T} \wedge \gamma_1^\mathcal{S} \stackrel{q}{\sim} \gamma_1^\mathcal{T})$$

Thus, the command mapping preserves the state mapping, property (ii) for the HMG+ correct simulation.

These properties satisfy the definition for HMG+ correct simulation, and hence  $\mathcal{T}$  admits an HMG+ simulation of  $\mathcal{S}$  ( $\mathcal{T} \stackrel{\text{sim}}{\underset{\text{HMG}+}{\sim}} \mathcal{S}$ ). ■

**Theorem 9**  $\text{HMG}+ \doteq \{\text{SCq}, \text{QDt}, \text{R} \rightarrow\}$ ; that is, the HMG+ parameterized expressiveness simulation (correctness only) decomposes to query correspondence, theory-dependent query, and forward reachability.

*Proof:* By Lemma 7, if  $\mathcal{T} \stackrel{\text{sim}}{\underset{\text{HMG}+}{\sim}} \mathcal{S}$ , then  $\mathcal{S} \leq^\mathcal{P} \mathcal{T}$ . By Lemma 8, if  $\mathcal{S} \leq^\mathcal{P} \mathcal{T}$ , then  $\mathcal{T} \stackrel{\text{sim}}{\underset{\text{HMG}+}{\sim}} \mathcal{S}$ . Thus,  $\mathcal{S} \leq^\mathcal{P} \mathcal{T}$  if and only if  $\mathcal{T} \stackrel{\text{sim}}{\underset{\text{HMG}+}{\sim}} \mathcal{S}$ , and thus the HMG+ simulation decomposes to  $\{\text{SCq}, \text{QDt}, \text{R} \rightarrow\}$ . ■

### C. AC-Preserving HMG+ Parameterized Expressiveness

For the purposes of this proof, let the set of simulation properties  $\mathcal{P} = \{\text{SCq}, \text{QDt}, \text{R} \rightarrow, \text{QPa}\}$ .

**Lemma 10** Given access control systems  $\mathcal{S}$ ,  $\mathcal{T}$ , and  $\mathcal{U}$ ,

$$\mathcal{T} \stackrel{\text{sim}}{\underset{\text{HMG}+a}{\sim}} \mathcal{S} \wedge \mathcal{S} \stackrel{\text{sim}}{\underset{\mathcal{P}}{\sim}} \mathcal{U} \Rightarrow \mathcal{T} \stackrel{\text{sim}}{\underset{\mathcal{P}}{\sim}} \mathcal{U}$$

That is, if  $\mathcal{T}$  admits an HMG+ simulation with AC-preservation of  $\mathcal{S}$ , and  $\mathcal{S}$  admits a simulation of  $\mathcal{U}$  with properties

$\{\text{SCq}, \text{QDt}, \text{R} \rightarrow, \text{QPa}\}$ , then  $\mathcal{T}$  admits a simulation of  $\mathcal{U}$  with properties  $\{\text{SCq}, \text{QDt}, \text{R} \rightarrow, \text{QPa}\}$ .

*Proof:* The proof of Lemma 7 proves all properties besides QPa with no change. Thus, we now show QPa.

Choose an arbitrary request  $r \in \mathcal{R}^\mathcal{U}$  and states  $\gamma^\mathcal{S} \in \Gamma^\mathcal{S}, \gamma^\mathcal{T} \in \Gamma^\mathcal{T}$ . Since  $\mathcal{S} \stackrel{\text{sim}}{\underset{\mathcal{P}}{\sim}} \mathcal{U}$ ,

$$\sigma_Q(r, \gamma^\mathcal{S}) \equiv \gamma^\mathcal{S} \vdash r$$

Since  $\mathcal{T} \stackrel{\text{sim}}{\underset{\text{HMG}+a}{\sim}} \mathcal{S}$ ,

$$\forall r^\mathcal{S} \in \mathcal{R}^\mathcal{S} : \sigma_Q(r^\mathcal{S}, \gamma^\mathcal{T}) \equiv \gamma^\mathcal{T} \vdash r^\mathcal{S}$$

Since  $\mathcal{S} \stackrel{\text{sim}}{\underset{\mathcal{P}}{\sim}} \mathcal{U}$ , by SCq,  $r \in \mathcal{R}^\mathcal{S}$ . Thus,

$$\sigma_Q(r, \gamma^\mathcal{T}) \equiv \gamma^\mathcal{T} \vdash r$$

Hence,  $\mathcal{T} \stackrel{\text{sim}}{\underset{\{\text{SCq}, \text{QDt}, \text{R} \rightarrow, \text{QPa}\}}{\sim}} \mathcal{U}$ . ■

**Lemma 11** Given access control systems  $\mathcal{S}$  and  $\mathcal{T}$  and simulation properties  $\mathcal{P} = \{\text{SCq}, \text{QDt}, \text{R} \rightarrow, \text{QPa}\}$ ,  $\mathcal{S} \leq^\mathcal{P} \mathcal{T}$ . That is, if  $\mathcal{T}$  is at least as expressive as  $\mathcal{S}$  with respect to properties  $\mathcal{P}$ , then  $\mathcal{T}$  admits an HMG+ simulation with AC-preservation of  $\mathcal{S}$ .

*Proof:* Let  $\mathcal{S}$  and  $\mathcal{T}$  be arbitrary access control systems such that  $\mathcal{S} \leq^\mathcal{P} \mathcal{T}$ . Since  $\mathcal{S} \leq^\mathcal{P} \mathcal{T}$ , for any access control system  $\mathcal{U}$ , if  $\mathcal{S} \stackrel{\text{sim}}{\underset{\mathcal{P}}{\sim}} \mathcal{U}$ , then  $\mathcal{T} \stackrel{\text{sim}}{\underset{\mathcal{P}}{\sim}} \mathcal{U}$ .

Since  $\mathcal{S}$  can trivially simulate itself,  $\mathcal{S} \stackrel{\text{sim}}{\underset{\mathcal{P}}{\sim}} \mathcal{S}$ , and thus  $\mathcal{T} \stackrel{\text{sim}}{\underset{\mathcal{P}}{\sim}} \mathcal{S}$ .

The proof of Lemma 8 proves all properties besides AC-preservation. By QPa,

$$\forall r^\mathcal{S} \in \mathcal{R}^\mathcal{S} : \sigma_Q(r^\mathcal{S}, \gamma^\mathcal{T}) \equiv \gamma^\mathcal{T} \vdash r^\mathcal{S}$$

This satisfies the definition of AC-preservation, and hence  $\mathcal{T}$  admits an HMG+ simulation with AC-preservation of  $\mathcal{S}$  ( $\mathcal{T} \stackrel{\text{sim}}{\underset{\text{HMG}+a}{\sim}} \mathcal{S}$ ). ■

**Theorem 12**  $\text{HMG}+a \doteq \{\text{SCq}, \text{QDt}, \text{R} \rightarrow, \text{QPa}\}$ ; that is, HMG+ parameterized expressiveness with AC-preservation decomposes to query correspondence, theory-dependent query, forward reachability, and authorization preservation.

*Proof:* By Lemma 10, if  $\mathcal{T} \stackrel{\text{sim}}{\underset{\text{HMG}+a}{\sim}} \mathcal{S}$ , then  $\mathcal{S} \leq^\mathcal{P} \mathcal{T}$ . By Lemma 11, if  $\mathcal{S} \leq^\mathcal{P} \mathcal{T}$ , then  $\mathcal{T} \stackrel{\text{sim}}{\underset{\text{HMG}+a}{\sim}} \mathcal{S}$ . Thus,  $\mathcal{S} \leq^\mathcal{P} \mathcal{T}$  if and only if  $\mathcal{T} \stackrel{\text{sim}}{\underset{\text{HMG}+a}{\sim}} \mathcal{S}$ , and thus the HMG+ simulation with AC-preservation decomposes to  $\{\text{SCq}, \text{QDt}, \text{R} \rightarrow, \text{QPa}\}$ . ■

### D. Monotonic HMG+ Parameterized Expressiveness

For the purposes of this proof, let the set of simulation properties  $\mathcal{P} = \{\text{SCq}, \text{QDt}, \text{R} \rightarrow, \text{CTa}\}$ .

**Lemma 13** Given access control systems  $\mathcal{S}$ ,  $\mathcal{T}$ , and  $\mathcal{U}$ ,

$$\mathcal{T} \stackrel{\text{sim}}{\underset{\text{HMG}+s}{\sim}} \mathcal{S} \wedge \mathcal{S} \stackrel{\text{sim}}{\underset{\mathcal{P}}{\sim}} \mathcal{U} \Rightarrow \mathcal{T} \stackrel{\text{sim}}{\underset{\mathcal{P}}{\sim}} \mathcal{U}$$

That is, if  $\mathcal{T}$  admits a monotonic HMG+ simulation of  $\mathcal{S}$ , and  $\mathcal{S}$  admits a simulation of  $\mathcal{U}$  with properties

$\{\text{SCq}, \text{QDt}, \text{R} \rightarrow, \text{CTa}\}$ , then  $\mathcal{T}$  admits a simulation of  $\mathcal{U}$  with properties  $\{\text{SCq}, \text{QDt}, \text{R} \rightarrow, \text{CTa}\}$ .

*Proof:* The proof of Lemma 7 proves all properties besides CTa with no change. Thus, we now show CTa.

Choose an arbitrary command  $\psi^{\mathcal{U}} \in \Psi^{\mathcal{U}}$  and state  $\gamma_0^{\mathcal{S}} \in \Gamma^{\mathcal{S}}$ . Let:

- $\langle \psi_1^{\mathcal{S}}, \dots, \psi_n^{\mathcal{S}} \rangle = \sigma_{\Psi}(\psi^{\mathcal{U}}, \gamma_0^{\mathcal{S}})$
- $\gamma_i^{\mathcal{S}} = \text{terminal}(\gamma_0^{\mathcal{S}}, \psi_1^{\mathcal{S}} \circ \dots \circ \psi_i^{\mathcal{S}})$

Since  $\mathcal{S} \xrightarrow{\mathcal{P}} \text{sim} \mathcal{U}$ , by CTa, this sequence of states  $\gamma_i^{\mathcal{S}}$  is monotonic.

Furthermore, let:

- $\gamma_{0,0}^{\mathcal{T}} = \sigma_{\Gamma}(\gamma_0^{\mathcal{S}})$
- $\langle \psi_{i,1}^{\mathcal{T}}, \dots, \psi_{i,m}^{\mathcal{T}} \rangle = \sigma_{\Psi}(\psi_i^{\mathcal{S}}, \gamma_{i-1,0}^{\mathcal{T}})$
- $\gamma_{i,0}^{\mathcal{T}} = \text{terminal}(\gamma_{i-1,0}^{\mathcal{T}}, \psi_{i,1}^{\mathcal{T}} \circ \dots \circ \psi_{i,m}^{\mathcal{T}})$
- $\gamma_{i,j}^{\mathcal{T}} = \text{terminal}(\gamma_{i,0}^{\mathcal{T}}, \psi_{i+1,1}^{\mathcal{T}} \circ \dots \circ \psi_{i+1,j}^{\mathcal{T}})$

Put simply,  $\psi^{\mathcal{U}}$  is simulated in  $\mathcal{T}$  by the following sequence of commands:

$$\psi_{1,1}^{\mathcal{T}}, \dots, \psi_{1,m}^{\mathcal{T}}, \psi_{2,1}^{\mathcal{T}}, \dots, \psi_{n,m}^{\mathcal{T}}$$

thus passing through the following trace of states:

$$\gamma_{0,0}^{\mathcal{T}}, \dots, \gamma_{0,m-1}^{\mathcal{T}}, \gamma_{1,0}^{\mathcal{T}}, \dots, \gamma_{n-1,m-1}^{\mathcal{T}}, \gamma_{n,0}^{\mathcal{T}}$$

Since  $\mathcal{T} \xrightarrow{\text{HMG}+s} \text{sim} \mathcal{S}$ , each subsequence  $\gamma_{i,0}^{\mathcal{T}}, \dots, \gamma_{i+1,0}^{\mathcal{T}}$  is monotonic.

Since  $\mathcal{S} \xrightarrow{\mathcal{P}} \text{sim} \mathcal{U}$ , by CTa and SCq, the full sequence must also be monotonic.

Hence,  $\mathcal{T} \xrightarrow{\{\text{SCq}, \text{QDt}, \text{R} \rightarrow, \text{CTa}\}} \text{sim} \mathcal{U}$ . ■

**Lemma 14** Given access control systems  $\mathcal{S}$  and  $\mathcal{T}$  and simulation properties  $\mathcal{P} = \{\text{SCq}, \text{QDt}, \text{R} \rightarrow, \text{CTa}\}$ ,  $\mathcal{S} \leq^{\mathcal{P}} \mathcal{T} \Rightarrow \mathcal{T} \xrightarrow{\text{HMG}+s} \text{sim} \mathcal{S}$ . That is, if  $\mathcal{T}$  is at least as expressive as  $\mathcal{S}$  with respect to properties  $\mathcal{P}$ , then  $\mathcal{T}$  admits a monotonic HMG+ simulation of  $\mathcal{S}$ .

*Proof:* Let  $\mathcal{S}$  and  $\mathcal{T}$  be arbitrary access control systems such that  $\mathcal{S} \leq^{\mathcal{P}} \mathcal{T}$ . Since  $\mathcal{S} \leq^{\mathcal{P}} \mathcal{T}$ , for any access control system  $\mathcal{U}$ , if  $\mathcal{S} \xrightarrow{\mathcal{P}} \text{sim} \mathcal{U}$ , then  $\mathcal{T} \xrightarrow{\mathcal{P}} \text{sim} \mathcal{U}$ .

Since  $\mathcal{S}$  can trivially simulate itself,  $\mathcal{S} \xrightarrow{\mathcal{P}} \text{sim} \mathcal{S}$ , and thus  $\mathcal{T} \xrightarrow{\mathcal{P}} \text{sim} \mathcal{S}$ .

The proof of Lemma 8 proves all properties besides monotonicity. Since CTa satisfies the definition of monotonicity, we thus have that  $\mathcal{T}$  admits a monotonic HMG+ simulation of  $\mathcal{S}$  ( $\mathcal{T} \xrightarrow{\text{HMG}+s} \text{sim} \mathcal{S}$ ). ■

**Theorem 15**  $\text{HMG}+s \doteq \{\text{SCq}, \text{QDt}, \text{R} \rightarrow, \text{CTa}\}$ ; that is, HMG+ parameterized expressiveness with monotonicity decomposes to query correspondence, theory-dependent query, forward reachability, and access monotonicity.

*Proof:* By Lemma 13, if  $\mathcal{T} \xrightarrow{\text{HMG}+s} \text{sim} \mathcal{S}$ , then  $\mathcal{S} \leq^{\mathcal{P}} \mathcal{T}$ . By Lemma 14, if  $\mathcal{S} \leq^{\mathcal{P}} \mathcal{T}$ , then  $\mathcal{T} \xrightarrow{\text{HMG}+s} \text{sim} \mathcal{S}$ . Thus,  $\mathcal{S} \leq^{\mathcal{P}} \mathcal{T}$  if and only if  $\mathcal{T} \xrightarrow{\text{HMG}+s} \text{sim} \mathcal{S}$ , and thus the monotonic HMG+ simulation decomposes to  $\{\text{SCq}, \text{QDt}, \text{R} \rightarrow, \text{CTa}\}$ . ■

**E. Admin-Preserving HMG+ Parameterized Expressiveness**

For the purposes of this proof, let the set of simulation properties  $\mathcal{P} = \{\text{SCq}, \text{QDt}, \text{R} \rightarrow, \text{CAa}\}$ .

**Lemma 16** Given access control systems  $\mathcal{S}$ ,  $\mathcal{T}$ , and  $\mathcal{U}$ ,

$$\mathcal{T} \xrightarrow{\text{HMG}+p} \text{sim} \mathcal{S} \wedge \mathcal{S} \xrightarrow{\mathcal{P}} \text{sim} \mathcal{U} \Rightarrow \mathcal{T} \xrightarrow{\mathcal{P}} \text{sim} \mathcal{U}$$

That is, if  $\mathcal{T}$  admits an admin-preserving HMG+ simulation of  $\mathcal{S}$ , and  $\mathcal{S}$  admits a simulation of  $\mathcal{U}$  with properties  $\{\text{SCq}, \text{QDt}, \text{R} \rightarrow, \text{CAa}\}$ , then  $\mathcal{T}$  admits a simulation of  $\mathcal{U}$  with properties  $\{\text{SCq}, \text{QDt}, \text{R} \rightarrow, \text{CAa}\}$ .

*Proof:* The proof of Lemma 7 proves all properties besides CAa with no change. Thus, we now show CAa.

Consider arbitrary command  $\psi^{\mathcal{U}} \in \Psi^{\mathcal{U}}$  and state  $\gamma^{\mathcal{S}} \in \Gamma^{\mathcal{S}}$ . Since  $\mathcal{S} \xrightarrow{\mathcal{P}} \text{sim} \mathcal{U}$ , by CAa,

$$\forall \psi^{\mathcal{S}} \in \sigma_{\Psi}(\psi^{\mathcal{U}}, \gamma^{\mathcal{S}}), \alpha(\psi^{\mathcal{S}}) \in A \Rightarrow \alpha(\psi^{\mathcal{U}}) \in A$$

Consider state  $\gamma^{\mathcal{T}} \in \Gamma^{\mathcal{T}}$ . Since  $\mathcal{T} \xrightarrow{\text{HMG}+p} \text{sim} \mathcal{S}$ , by admin-preservation,

$$\forall \psi^{\mathcal{T}} \in \sigma_{\Psi}(\psi^{\mathcal{S}}, \gamma^{\mathcal{T}}), \alpha(\psi^{\mathcal{T}}) \in A \Rightarrow \alpha(\psi^{\mathcal{S}}) \in A$$

Thus,

$$\forall \psi^{\mathcal{T}} \in \sigma_{\Psi}(\psi^{\mathcal{U}}, \gamma^{\mathcal{T}}), \alpha(\psi^{\mathcal{T}}) \in A \Rightarrow \alpha(\psi^{\mathcal{U}}) \in A$$

Hence,  $\mathcal{T} \xrightarrow{\{\text{SCq}, \text{QDt}, \text{R} \rightarrow, \text{CAa}\}} \text{sim} \mathcal{U}$ . ■

**Lemma 17** Given access control systems  $\mathcal{S}$  and  $\mathcal{T}$  and simulation properties  $\mathcal{P} = \{\text{SCq}, \text{QDt}, \text{R} \rightarrow, \text{CAa}\}$ ,  $\mathcal{S} \leq^{\mathcal{P}} \mathcal{T} \Rightarrow \mathcal{T} \xrightarrow{\text{HMG}+p} \text{sim} \mathcal{S}$ . That is, if  $\mathcal{T}$  is at least as expressive as  $\mathcal{S}$  with respect to properties  $\mathcal{P}$ , then  $\mathcal{T}$  admits an admin-preserving HMG+ simulation of  $\mathcal{S}$ .

*Proof:* Let  $\mathcal{S}$  and  $\mathcal{T}$  be arbitrary access control systems such that  $\mathcal{S} \leq^{\mathcal{P}} \mathcal{T}$ . Since  $\mathcal{S} \leq^{\mathcal{P}} \mathcal{T}$ , for any access control system  $\mathcal{U}$ , if  $\mathcal{S} \xrightarrow{\mathcal{P}} \text{sim} \mathcal{U}$ , then  $\mathcal{T} \xrightarrow{\mathcal{P}} \text{sim} \mathcal{U}$ .

Since  $\mathcal{S}$  can trivially simulate itself,  $\mathcal{S} \xrightarrow{\mathcal{P}} \text{sim} \mathcal{S}$ , and thus  $\mathcal{T} \xrightarrow{\mathcal{P}} \text{sim} \mathcal{S}$ .

The proof of Lemma 8 proves all properties besides admin-preservation. Since CAa satisfies the definition of admin-preservation, we thus have that  $\mathcal{T}$  admits an admin-preserving HMG+ simulation of  $\mathcal{S}$  ( $\mathcal{T} \xrightarrow{\text{HMG}+p} \text{sim} \mathcal{S}$ ). ■

**Theorem 18**  $\text{HMG}+p \doteq \{\text{SCq}, \text{QDt}, \text{R} \rightarrow, \text{CAa}\}$ ; that is, HMG+ parameterized expressiveness with admin-preservation decomposes to query correspondence, theory-dependent query, forward reachability, and administration preservation.

*Proof:* By Lemma 16, if  $\mathcal{T} \xrightarrow{\text{HMG}+p} \text{sim} \mathcal{S}$ , then  $\mathcal{S} \leq^{\mathcal{P}} \mathcal{T}$ . By Lemma 17, if  $\mathcal{S} \leq^{\mathcal{P}} \mathcal{T}$ , then  $\mathcal{T} \xrightarrow{\text{HMG}+p} \text{sim} \mathcal{S}$ . Thus,  $\mathcal{S} \leq^{\mathcal{P}} \mathcal{T}$  if and only if  $\mathcal{T} \xrightarrow{\text{HMG}+p} \text{sim} \mathcal{S}$ , and thus the admin-preserving HMG+ simulation decomposes to  $\{\text{SCq}, \text{QDt}, \text{R} \rightarrow, \text{CAa}\}$ . ■

## F. SMG Simulation

For the purposes of this proof, let the set of simulation properties  $\mathcal{P} = \{\text{SCa}, \text{R} \rightarrow\}$ .

**Lemma 19** Given access control systems  $\mathcal{S}$ ,  $\mathcal{T}$ , and  $\mathcal{U}$ ,

$$\mathcal{T} \stackrel{\text{sim}}{\underset{\text{SMG}}{\sim}} \mathcal{S} \wedge \mathcal{S} \stackrel{\text{sim}}{\underset{\mathcal{P}}{\sim}} \mathcal{U} \Rightarrow \mathcal{T} \stackrel{\text{sim}}{\underset{\mathcal{P}}{\sim}} \mathcal{U}$$

That is, if  $\mathcal{T}$  admits an SMG simulation of  $\mathcal{S}$ , and  $\mathcal{S}$  admits a simulation of  $\mathcal{U}$  with properties  $\{\text{SCa}, \text{R} \rightarrow\}$ , then  $\mathcal{T}$  admits a simulation of  $\mathcal{U}$  with properties  $\{\text{SCa}, \text{R} \rightarrow\}$ .

*Proof:* Let  $\mathcal{S}$ ,  $\mathcal{T}$ , and  $\mathcal{U}$  be arbitrary access control systems such that  $\mathcal{T} \stackrel{\text{sim}}{\underset{\text{SMG}}{\sim}} \mathcal{S}$  and  $\mathcal{S} \stackrel{\text{sim}}{\underset{\mathcal{P}}{\sim}} \mathcal{U}$ . To prove Lemma 19, we must then show that  $\mathcal{T} \stackrel{\text{sim}}{\underset{\mathcal{P}}{\sim}} \mathcal{U}$ .

Choose an arbitrary state  $\gamma_0^\mathcal{U} \in \Gamma^\mathcal{U}$  and command  $\psi^\mathcal{U} \in \Psi^\mathcal{U}$ , and let  $\text{next}(\gamma_0^\mathcal{U}, \psi^\mathcal{U}) = \gamma_1^\mathcal{U}$ . Let  $\gamma_0^\mathcal{S} \in \Gamma^\mathcal{S}$  such that  $\gamma_0^\mathcal{U} \stackrel{a}{\sim} \gamma_0^\mathcal{S}$ . Since  $\mathcal{S} \stackrel{\text{sim}}{\underset{\mathcal{P}}{\sim}} \mathcal{U}$ ,

$$\exists \gamma_1^\mathcal{S} \in \Gamma^\mathcal{S}. (\text{terminal}(\gamma_0^\mathcal{S}, \sigma_\Psi(\psi^\mathcal{U}, \gamma_0^\mathcal{S})) = \gamma_1^\mathcal{S} \wedge \gamma_1^\mathcal{U} \stackrel{a}{\sim} \gamma_1^\mathcal{S})$$

Let  $\gamma_0^\mathcal{T} \in \Gamma^\mathcal{T}$  such that  $\gamma_0^\mathcal{S} \stackrel{a}{\sim} \gamma_0^\mathcal{T}$ . Since  $\mathcal{T} \stackrel{\text{sim}}{\underset{\text{SMG}}{\sim}} \mathcal{S}$ , request  $r$  can be granted in  $\mathcal{S}$  if and only if  $\sigma(r)$  can be granted in  $\mathcal{T}$ . More concretely,

$$\exists \gamma_1^\mathcal{T} \in \Gamma^\mathcal{T}. (\gamma_0^\mathcal{T} \stackrel{*}{\mapsto} \gamma_1^\mathcal{T} \wedge \gamma_1^\mathcal{S} \stackrel{a}{\sim} \gamma_1^\mathcal{T})$$

Thus, there exists a sequence of  $\mathcal{T}$  commands  $\Psi_0^\mathcal{T}$  such that  $\text{terminal}(\gamma_0^\mathcal{T}, \Psi_0^\mathcal{T}) = \gamma_1^\mathcal{T}$ . Define  $\sigma_\Psi : \Psi^\mathcal{U} \times \Gamma^\mathcal{T} \rightarrow (\Psi^\mathcal{T})^*$  such that it returns  $\Psi_0^\mathcal{T}$  for  $\gamma_0^\mathcal{T}, \psi^\mathcal{U}$ . This is formed by concatenating a sequence of sequences of commands: for each command  $\psi_i^\mathcal{S}$  that  $\mathcal{S}$  needs to execute to simulate  $\psi^\mathcal{U}$ , concatenate the commands that  $\mathcal{T}$  needs to execute to simulate  $\psi_i^\mathcal{S}$ .

Then, given  $\gamma_0^\mathcal{U}, \gamma_1^\mathcal{U} \in \Gamma^\mathcal{U}, \gamma_0^\mathcal{T} \in \Gamma^\mathcal{T}, \psi^\mathcal{U} \in \Psi^\mathcal{U}$  such that  $\text{next}(\gamma_0^\mathcal{U}, \psi^\mathcal{U}) = \gamma_1^\mathcal{U}$ , and  $\gamma_0^\mathcal{U} \stackrel{a}{\sim} \gamma_0^\mathcal{T}$ ,

$$\exists \gamma_1^\mathcal{T} \in \Gamma^\mathcal{T}. (\text{terminal}(\gamma_0^\mathcal{T}, \sigma_\Psi(\psi^\mathcal{U}, \gamma_0^\mathcal{T})) = \gamma_1^\mathcal{T} \wedge \gamma_1^\mathcal{S} \stackrel{a}{\sim} \gamma_1^\mathcal{T})$$

Hence,  $\mathcal{T} \stackrel{\text{sim}}{\underset{\{\text{SCa}, \text{R} \rightarrow\}}{\sim}} \mathcal{U}$ . ■

**Lemma 20** Given access control systems  $\mathcal{S}$  and  $\mathcal{T}$  and simulation properties  $\mathcal{P} = \{\text{SCa}, \text{R} \rightarrow\}$ ,  $\mathcal{S} \leq^\mathcal{P} \mathcal{T} \Rightarrow \mathcal{T} \stackrel{\text{sim}}{\underset{\text{SMG}}{\sim}} \mathcal{S}$ . That is, if  $\mathcal{T}$  is at least as expressive as  $\mathcal{S}$  with respect to properties  $\mathcal{P}$ , then  $\mathcal{T}$  admits an SMG simulation of  $\mathcal{S}$ .

*Proof:* Let  $\mathcal{S}$  and  $\mathcal{T}$  be arbitrary access control systems such that  $\mathcal{S} \leq^\mathcal{P} \mathcal{T}$ . Since  $\mathcal{S} \leq^\mathcal{P} \mathcal{T}$ , for any access control system  $\mathcal{U}$ , if  $\mathcal{S} \stackrel{\text{sim}}{\underset{\mathcal{P}}{\sim}} \mathcal{U}$ , then  $\mathcal{T} \stackrel{\text{sim}}{\underset{\mathcal{P}}{\sim}} \mathcal{U}$ .

Since  $\mathcal{S}$  can trivially simulate itself,  $\mathcal{S} \stackrel{\text{sim}}{\underset{\mathcal{P}}{\sim}} \mathcal{S}$ , and thus  $\mathcal{T} \stackrel{\text{sim}}{\underset{\mathcal{P}}{\sim}} \mathcal{S}$ .

By SCa and QPa, request  $r$  can be granted in  $\mathcal{S}$  if and only if  $\sigma(r)$  can be granted in  $\mathcal{T}$ . This satisfies the definition of the SMG simulation, and hence  $\mathcal{T}$  admits an SMG simulation of  $\mathcal{S}$  ( $\mathcal{T} \stackrel{\text{sim}}{\underset{\text{SMG}}{\sim}} \mathcal{S}$ ). ■

**Theorem 21**  $\text{SMG} \doteq \{\text{SCa}, \text{R} \rightarrow\}$ ; that is, the SMG simulation decomposes to authorization correspondence and forward reachability.

*Proof:* By Lemma 19, if  $\mathcal{T} \stackrel{\text{sim}}{\underset{\text{SMG}}{\sim}} \mathcal{S}$ , then  $\mathcal{S} \leq^\mathcal{P} \mathcal{T}$ . By Lemma 20, if  $\mathcal{S} \leq^\mathcal{P} \mathcal{T}$ , then  $\mathcal{T} \stackrel{\text{sim}}{\underset{\text{SMG}}{\sim}} \mathcal{S}$ . Thus,  $\mathcal{S} \leq^\mathcal{P} \mathcal{T}$  if and only if  $\mathcal{T} \stackrel{\text{sim}}{\underset{\text{SMG}}{\sim}} \mathcal{S}$ , and thus the SMG simulation decomposes to  $\{\text{SCa}, \text{R} \rightarrow\}$ . ■

## G. Ganta Simulation

For the purposes of this proof, let the set of simulation properties  $\mathcal{P} = \{\text{SCa}, \text{QPa}, \text{CTs}, \text{R} \leftrightarrow\}$ .

**Lemma 22** Given access control systems  $\mathcal{S}$ ,  $\mathcal{T}$ , and  $\mathcal{U}$ ,

$$\mathcal{T} \stackrel{\text{sim}}{\underset{\text{Ganta}}{\sim}} \mathcal{S} \wedge \mathcal{S} \stackrel{\text{sim}}{\underset{\mathcal{P}}{\sim}} \mathcal{U} \Rightarrow \mathcal{T} \stackrel{\text{sim}}{\underset{\mathcal{P}}{\sim}} \mathcal{U}$$

That is, if  $\mathcal{T}$  admits a Ganta simulation of  $\mathcal{S}$ , and  $\mathcal{S}$  admits a simulation of  $\mathcal{U}$  with properties  $\{\text{SCa}, \text{QPa}, \text{CTs}, \text{R} \leftrightarrow\}$ , then  $\mathcal{T}$  admits a simulation of  $\mathcal{U}$  with properties  $\{\text{SCa}, \text{QPa}, \text{CTs}, \text{R} \leftrightarrow\}$ .

*Proof:* Let  $\mathcal{S}$ ,  $\mathcal{T}$ , and  $\mathcal{U}$  be arbitrary access control systems such that  $\mathcal{T} \stackrel{\text{sim}}{\underset{\text{Ganta}}{\sim}} \mathcal{S}$  and  $\mathcal{S} \stackrel{\text{sim}}{\underset{\mathcal{P}}{\sim}} \mathcal{U}$ . To prove Lemma 22, we must then show that  $\mathcal{T} \stackrel{\text{sim}}{\underset{\mathcal{P}}{\sim}} \mathcal{U}$ .

Choose an arbitrary state  $\gamma_0^\mathcal{U} \in \Gamma^\mathcal{U}$  and command  $\psi^\mathcal{U} \in \Psi^\mathcal{U}$ , and let  $\text{next}(\gamma_0^\mathcal{U}, \psi^\mathcal{U}) = \gamma_1^\mathcal{U}$ . Let  $\gamma_0^\mathcal{S} \in \Gamma^\mathcal{S}$  such that  $\gamma_0^\mathcal{U} \stackrel{a}{\sim} \gamma_0^\mathcal{S}$ . Since  $\mathcal{S} \stackrel{\text{sim}}{\underset{\mathcal{P}}{\sim}} \mathcal{U}$ ,

$$\exists \gamma_1^\mathcal{S} \in \Gamma^\mathcal{S}. (\text{terminal}(\gamma_0^\mathcal{S}, \sigma_\Psi(\psi^\mathcal{U}, \gamma_0^\mathcal{S})) = \gamma_1^\mathcal{S} \wedge \gamma_1^\mathcal{U} \stackrel{a}{\sim} \gamma_1^\mathcal{S})$$

Let  $\gamma_0^\mathcal{T} \in \Gamma^\mathcal{T}$  such that  $\gamma_0^\mathcal{S} \stackrel{a}{\sim} \gamma_0^\mathcal{T}$ . Since  $\mathcal{T} \stackrel{\text{sim}}{\underset{\text{Ganta}}{\sim}} \mathcal{S}$ , by Property 2 there must exist an equivalent history to  $\gamma_0^\mathcal{S} \stackrel{*}{\mapsto} \gamma_1^\mathcal{S}$  in  $\mathcal{T}$  with an access-correspondent completion state. Thus,

$$\exists \gamma_1^\mathcal{T} \in \Gamma^\mathcal{T}. (\gamma_0^\mathcal{T} \stackrel{*}{\mapsto} \gamma_1^\mathcal{T} \wedge \gamma_1^\mathcal{S} \stackrel{a}{\sim} \gamma_1^\mathcal{T})$$

Thus, there exists a sequence of  $\mathcal{T}$  commands  $\Psi_0^\mathcal{T}$  such that  $\text{terminal}(\gamma_0^\mathcal{T}, \Psi_0^\mathcal{T}) = \gamma_1^\mathcal{T}$ . Define  $\sigma_\Psi : \Psi^\mathcal{U} \times \Gamma^\mathcal{T} \rightarrow (\Psi^\mathcal{T})^*$  such that it returns  $\Psi_0^\mathcal{T}$  for  $\gamma_0^\mathcal{T}, \psi^\mathcal{U}$ . This is formed by concatenating a sequence of sequences of commands: for each command  $\psi_i^\mathcal{S}$  that  $\mathcal{S}$  needs to execute to simulate  $\psi^\mathcal{U}$ , concatenate the commands that  $\mathcal{T}$  needs to execute to simulate  $\psi_i^\mathcal{S}$ .

Then, given  $\gamma_0^\mathcal{U}, \gamma_1^\mathcal{U} \in \Gamma^\mathcal{U}, \gamma_0^\mathcal{T} \in \Gamma^\mathcal{T}, \psi^\mathcal{U} \in \Psi^\mathcal{U}$  such that  $\text{next}(\gamma_0^\mathcal{U}, \psi^\mathcal{U}) = \gamma_1^\mathcal{U}$ , and  $\gamma_0^\mathcal{U} \stackrel{a}{\sim} \gamma_0^\mathcal{T}$ ,

$$\exists \gamma_1^\mathcal{T} \in \Gamma^\mathcal{T}. (\text{terminal}(\gamma_0^\mathcal{T}, \sigma_\Psi(\psi^\mathcal{U}, \gamma_0^\mathcal{T})) = \gamma_1^\mathcal{T} \wedge \gamma_1^\mathcal{S} \stackrel{a}{\sim} \gamma_1^\mathcal{T})$$

Hence,  $\mathcal{T} \stackrel{\text{sim}}{\underset{\{\text{SCa}, \text{R} \rightarrow\}}{\sim}} \mathcal{U}$ . Next we show  $\text{R} \leftrightarrow$ .

Choose some arbitrary states  $\gamma_0^\mathcal{T}, \gamma_1^\mathcal{T} \in \Gamma^\mathcal{T}$  such that  $\gamma_0^\mathcal{T} \mapsto \gamma_1^\mathcal{T}$ . Let  $\gamma_0^\mathcal{S} \in \Gamma^\mathcal{S}$  such that  $\gamma_0^\mathcal{T} \stackrel{a}{\sim} \gamma_0^\mathcal{S}$ . Since  $\mathcal{T} \stackrel{\text{sim}}{\underset{\text{Ganta}}{\sim}} \mathcal{S}$ , by Property 3 there must exist an equivalent history to  $\gamma_0^\mathcal{T} \mapsto \gamma_1^\mathcal{T}$  in  $\mathcal{S}$  with an access-correspondent completion state. Thus,

$$\exists \gamma_1^\mathcal{S} \in \Gamma^\mathcal{S}. (\gamma_0^\mathcal{S} \stackrel{*}{\mapsto} \gamma_1^\mathcal{S} \wedge \gamma_1^\mathcal{T} \stackrel{a}{\sim} \gamma_1^\mathcal{S})$$

Let  $\gamma_0^\mathcal{U} \in \Gamma^\mathcal{U}$  such that  $\gamma_0^\mathcal{S} \stackrel{a}{\sim} \gamma_0^\mathcal{U}$ . Since  $\mathcal{S} \stackrel{\text{sim}}{\underset{\mathcal{P}}{\sim}} \mathcal{U}$ ,

$$\exists \gamma_1^\mathcal{U} \in \Gamma^\mathcal{U}. (\gamma_0^\mathcal{U} \stackrel{*}{\mapsto} \gamma_1^\mathcal{U} \wedge \gamma_1^\mathcal{T} \stackrel{a}{\sim} \gamma_1^\mathcal{U})$$

Thus, given  $\gamma_0^T, \gamma_1^T \in \Gamma^T, \gamma_0^U \in \Gamma^U$  such that  $\gamma_0^T \mapsto \gamma_1^T$  and  $\gamma_0^U \stackrel{a}{\sim} \gamma_0^T$ ,

$$\exists \gamma_1^U \in \Gamma^U. (\gamma_0^U \mapsto \gamma_1^U \wedge \gamma_1^U \stackrel{a}{\sim} \gamma_1^T)$$

Hence,  $\mathcal{T} \stackrel{\text{sim}}{\underset{\{\text{SCa}, \text{R} \leftrightarrow\}}{\mathcal{U}}}$ . Next we show QPa.

Choose an arbitrary request  $r \in \mathcal{R}^U$  and states  $\gamma^S \in \Gamma^S, \gamma^T \in \Gamma^T$ . Since  $\mathcal{S} \stackrel{\text{sim}}{\underset{\mathcal{P}}{\mathcal{U}}}$ ,

$$\sigma_Q(r, \gamma^S) \equiv \gamma^S \vdash r$$

Since  $\mathcal{T} \stackrel{\text{sim}}{\underset{\text{Ganta}}{\mathcal{S}}}$ , by Property 6,

$$\forall r^S \in \mathcal{R}^S : \sigma_Q(r^S, \gamma^T) \equiv \gamma^T \vdash r^S$$

Since  $\mathcal{S} \stackrel{\text{sim}}{\underset{\mathcal{P}}{\mathcal{U}}}$ , by SCa,  $r \in \mathcal{R}^S$ . Thus,

$$\sigma_Q(r, \gamma^T) \equiv \gamma^T \vdash r$$

Hence,  $\mathcal{T} \stackrel{\text{sim}}{\underset{\{\text{SCa}, \text{QPa}, \text{R} \leftrightarrow\}}{\mathcal{U}}}$ . Next we show CTs.

Choose an arbitrary command  $\psi^U \in \Psi^U$  and state  $\gamma_0^S \in \Gamma^S$ . Let:

- $\langle \psi_1^S, \dots, \psi_n^S \rangle = \sigma_\Psi(\psi^U, \gamma_0^S)$
- $\gamma_i^S = \text{terminal}(\gamma_0^S, \psi_1^S \circ \dots \circ \psi_i^S)$

Since  $\mathcal{S} \stackrel{\text{sim}}{\underset{\mathcal{P}}{\mathcal{U}}}$ , by CTs,

$$\text{Allowed}(\gamma_i^S) \subseteq \text{Allowed}(\gamma_0^S) \vee \text{Allowed}(\gamma_i^S) \subseteq \text{Allowed}(\gamma_n^S)$$

for any  $\gamma_i^S$ .

Furthermore, let:

- $\gamma_{0,0}^T = \sigma_\Gamma(\gamma_0^S)$
- $\langle \psi_{i,1}^T, \dots, \psi_{i,m}^T \rangle = \sigma_\Psi(\psi_i^S, \gamma_{i-1,0}^T)$
- $\gamma_{i,0}^T = \text{terminal}(\gamma_{i-1,0}^T, \psi_{i,1}^T \circ \dots \circ \psi_{i,m}^T)$
- $\gamma_{i,j}^T = \text{terminal}(\gamma_{i,0}^T, \psi_{i+1,1}^T \circ \dots \circ \psi_{i+1,j}^T)$

Put simply,  $\psi^U$  is simulated in  $\mathcal{T}$  by the following sequence of commands:

$$\psi_{1,1}^T, \dots, \psi_{1,m}^T, \psi_{2,1}^T, \dots, \psi_{n,m}^T$$

thus passing through the following trace of states:

$$\gamma_{0,0}^T, \dots, \gamma_{0,m-1}^T, \gamma_{1,0}^T, \dots, \gamma_{n-1,m-1}^T, \gamma_{n,0}^T$$

Consider an arbitrary state  $\gamma_{i,j}^T$  in this trace. Since  $\mathcal{T} \stackrel{\text{sim}}{\underset{\text{Ganta}}{\mathcal{S}}}$ , by Property 6

$$\text{Allowed}(\gamma_{i,j}^T) \subseteq \text{Allowed}(\gamma_{i,0}^T) \vee \text{Allowed}(\gamma_{i,j}^T) \subseteq \text{Allowed}(\gamma_{i+1,0}^T)$$

Since  $\mathcal{S} \stackrel{\text{sim}}{\underset{\mathcal{P}}{\mathcal{U}}}$ , by CTs and SCa,

$$\text{Allowed}(\gamma_{i,0}^T) \subseteq \text{Allowed}(\gamma_{0,0}^T) \vee \text{Allowed}(\gamma_{i,0}^T) \subseteq \text{Allowed}(\gamma_{n,0}^T)$$

as well as for  $\gamma_{i+1,0}^T$  in place of  $\gamma_{i,0}^T$ .

Thus,

$$\text{Allowed}(\gamma_{i,j}^T) \subseteq \text{Allowed}(\gamma_{0,0}^T) \vee \text{Allowed}(\gamma_{i,j}^T) \subseteq \text{Allowed}(\gamma_{n,0}^T)$$

Hence,  $\mathcal{T} \stackrel{\text{sim}}{\underset{\{\text{SCa}, \text{QPa}, \text{CTs}, \text{R} \leftrightarrow\}}{\mathcal{U}}}$ . ■

**Lemma 23** Given access control systems  $\mathcal{S}$  and  $\mathcal{T}$  and simulation properties  $\mathcal{P} = \{\text{SCa}, \text{QPa}, \text{CTs}, \text{R} \leftrightarrow\}$ ,  $\mathcal{S} \leq^{\mathcal{P}} \mathcal{T} \Rightarrow$

$\mathcal{T} \stackrel{\text{sim}}{\underset{\text{Ganta}}{\mathcal{S}}}$ . That is, if  $\mathcal{T}$  is at least as expressive as  $\mathcal{S}$  with respect to properties  $\mathcal{P}$ , then  $\mathcal{T}$  admits a Ganta simulation of  $\mathcal{S}$ .

*Proof:* Let  $\mathcal{S}$  and  $\mathcal{T}$  be arbitrary access control systems such that  $\mathcal{S} \leq^{\mathcal{P}} \mathcal{T}$ . Since  $\mathcal{S} \leq^{\mathcal{P}} \mathcal{T}$ , for any access control system  $\mathcal{U}$ , if  $\mathcal{S} \stackrel{\text{sim}}{\underset{\mathcal{P}}{\mathcal{U}}}$ , then  $\mathcal{T} \stackrel{\text{sim}}{\underset{\mathcal{P}}{\mathcal{U}}}$ .

Since  $\mathcal{S}$  can trivially simulate itself,  $\mathcal{S} \stackrel{\text{sim}}{\underset{\mathcal{P}}{\mathcal{S}}}$ , and thus  $\mathcal{T} \stackrel{\text{sim}}{\underset{\mathcal{P}}{\mathcal{S}}}$ .

By SCa, QPa, and the definition of  $\sigma_\Psi$ , we have a Ganta scheme mapping from  $\mathcal{S}$  to  $\mathcal{T}$ , satisfying Ganta simulation Property 1.

By  $\text{R} \rightarrow$  and the definition of  $\sigma_\Psi$ , all histories of  $\mathcal{S}$  have equivalent serial histories of  $\mathcal{T}$ , satisfying Property 2.

By  $\text{R} \leftrightarrow$  and the definition of  $\sigma_\Psi$ , all incomplete histories of  $\mathcal{T}$  can be completed serially, and all complete histories of  $\mathcal{T}$  have equivalent histories of  $\mathcal{S}$ , satisfying Properties 3–5.

Finally, by QPa and CTs, completion states are access-correspondent, and intermediate states are non-contaminating, satisfying Property 6.

These properties define the Ganta simulation, and hence  $\mathcal{T}$  admits a Ganta simulation of  $\mathcal{S}$  ( $\mathcal{T} \stackrel{\text{sim}}{\underset{\text{Ganta}}{\mathcal{S}}}$ ). ■

**Theorem 24**  $\text{Ganta} \doteq \{\text{SCa}, \text{QPa}, \text{CTs}, \text{R} \leftrightarrow\}$ ; that is, the Ganta simulation decomposes to access correspondence, authorization preservation, anti-contamination, and bidirectional reachability.

*Proof:* By Lemma 22, if  $\mathcal{T} \stackrel{\text{sim}}{\underset{\text{Ganta}}{\mathcal{S}}}$ , then  $\mathcal{S} \leq^{\mathcal{P}} \mathcal{T}$ . By Lemma 23, if  $\mathcal{S} \leq^{\mathcal{P}} \mathcal{T}$ , then  $\mathcal{T} \stackrel{\text{sim}}{\underset{\text{Ganta}}{\mathcal{S}}}$ . Thus,  $\mathcal{S} \leq^{\mathcal{P}} \mathcal{T}$  if and only if  $\mathcal{T} \stackrel{\text{sim}}{\underset{\text{Ganta}}{\mathcal{S}}}$ , and thus the Ganta simulation decomposes to  $\{\text{SCa}, \text{QPa}, \text{CTs}, \text{R} \leftrightarrow\}$ . ■

#### H. CDM Weak Simulation

For the purposes of this proof, let the set of simulation properties  $\mathcal{P} = \{\text{SCa}, \text{QPa}, \text{CDi}, \text{R} \rightarrow\}$ .

**Lemma 25** Given access control systems  $\mathcal{S}$ ,  $\mathcal{T}$ , and  $\mathcal{U}$ ,

$$\mathcal{T} \stackrel{\text{sim}}{\underset{\text{CDMw}}{\mathcal{S}}} \wedge \mathcal{S} \stackrel{\text{sim}}{\underset{\mathcal{P}}{\mathcal{U}}} \Rightarrow \mathcal{T} \stackrel{\text{sim}}{\underset{\mathcal{P}}{\mathcal{U}}}$$

That is, if  $\mathcal{T}$  admits a CDM weak simulation of  $\mathcal{S}$ , and  $\mathcal{S}$  admits a simulation of  $\mathcal{U}$  with properties  $\{\text{SCa}, \text{QPa}, \text{CDi}, \text{R} \rightarrow\}$ , then  $\mathcal{T}$  admits a simulation of  $\mathcal{U}$  with properties  $\{\text{SCa}, \text{QPa}, \text{CDi}, \text{R} \rightarrow\}$ .

*Proof:* To prove this lemma, we let  $\mathcal{S}$ ,  $\mathcal{T}$ , and  $\mathcal{U}$  be access control systems such that  $\mathcal{T} \stackrel{\text{sim}}{\underset{\text{CDMw}}{\mathcal{S}}}$  and  $\mathcal{S} \stackrel{\text{sim}}{\underset{\mathcal{P}}{\mathcal{U}}}$  but are otherwise arbitrary, and we show that  $\mathcal{T} \stackrel{\text{sim}}{\underset{\mathcal{P}}{\mathcal{U}}}$ .

Choose an arbitrary state  $\gamma_0^U \in \Gamma^U$  and command  $\psi^U \in \Psi^U$ , and let  $\text{next}(\gamma_0^U, \psi^U) = \gamma_1^U$ . Let  $\gamma_0^S \in \Gamma^S$  such that  $\gamma_0^U \stackrel{a}{\sim} \gamma_0^S$ . Since  $\mathcal{S} \stackrel{\text{sim}}{\underset{\mathcal{P}}{\mathcal{U}}}$ ,

$$\exists \gamma_1^S \in \Gamma^S. (\text{terminal}(\gamma_0^S, \sigma_\Psi(\psi^U, \gamma_0^S)) = \gamma_1^S \wedge \gamma_1^U \stackrel{a}{\sim} \gamma_1^S)$$

Let  $\gamma_0^T \in \Gamma^T$  such that  $\gamma_0^S \stackrel{a}{\sim} \gamma_0^T$ . Since  $\mathcal{T} \stackrel{\text{sim}}{\underset{\text{CDMw}}{\mathcal{S}}}$ ,

$$\exists \gamma_1^T \in \Gamma^T. (\gamma_0^T \mapsto \gamma_1^T \wedge \gamma_1^S \stackrel{a}{\sim} \gamma_1^T)$$

Thus, there exists a sequence of  $\mathcal{T}$  commands  $\Psi_0^\mathcal{T}$  such that  $\text{terminal}(\gamma_0^\mathcal{T}, \Psi_0^\mathcal{T}) = \gamma_1^\mathcal{T}$ . Define  $\sigma_\Psi : \Psi^\mathcal{U} \times \Gamma^\mathcal{T} \rightarrow (\Psi^\mathcal{T})^*$  such that it returns  $\Psi_0^\mathcal{T}$  for  $\gamma_0^\mathcal{T}, \psi^\mathcal{U}$ .

Then, given  $\gamma_0^\mathcal{U}, \gamma_1^\mathcal{U} \in \Gamma^\mathcal{U}, \gamma_0^\mathcal{T} \in \Gamma^\mathcal{T}, \psi^\mathcal{U} \in \Psi^\mathcal{U}$  such that  $\text{next}(\gamma_0^\mathcal{U}, \psi^\mathcal{U}) = \gamma_1^\mathcal{U}$ , and  $\gamma_0^\mathcal{U} \stackrel{a}{\sim} \gamma_0^\mathcal{T}$ ,

$$\exists \gamma_1^\mathcal{T} \in \Gamma^\mathcal{T}. (\text{terminal}(\gamma_0^\mathcal{T}, \sigma_\Psi(\psi, \gamma_0^\mathcal{T})) = \gamma_1^\mathcal{T} \wedge \gamma_1^\mathcal{S} \stackrel{a}{\sim} \gamma_1^\mathcal{T})$$

Hence,  $\mathcal{T} \stackrel{\text{sim}}{\underset{\{\text{SCa}, \text{R} \rightarrow\}}{\mathcal{U}}}$ . Next, we show QPa.

Choose an arbitrary request  $r \in \mathcal{R}^\mathcal{U}$  and states  $\gamma^\mathcal{S} \in \Gamma^\mathcal{S}, \gamma^\mathcal{T} \in \Gamma^\mathcal{T}$ . Since  $\mathcal{S} \stackrel{\text{sim}}{\underset{\mathcal{P}}{\mathcal{U}}}$ ,

$$\sigma_Q(r, \gamma^\mathcal{S}) \equiv \gamma^\mathcal{S} \vdash r$$

Since  $\mathcal{T} \stackrel{\text{sim}}{\underset{\text{CDM}_w}{\mathcal{S}}}$ ,

$$\forall r^\mathcal{S} \in \mathcal{R}^\mathcal{S} : \sigma_Q(r^\mathcal{S}, \gamma^\mathcal{T}) \equiv \gamma^\mathcal{T} \vdash r^\mathcal{S}$$

Since  $\mathcal{S} \stackrel{\text{sim}}{\underset{\mathcal{P}}{\mathcal{U}}}$ , by SCa,  $r \in \mathcal{R}^\mathcal{S}$ . Thus,

$$\sigma_Q(r, \gamma^\mathcal{T}) \equiv \gamma^\mathcal{T} \vdash r$$

Hence,  $\mathcal{T} \stackrel{\text{sim}}{\underset{\{\text{SCa}, \text{QPa}, \text{R} \rightarrow\}}{\mathcal{U}}}$ . Next we show CDi.

Since  $\mathcal{S} \stackrel{\text{sim}}{\underset{\mathcal{P}}{\mathcal{U}}}$ , by CDi,  $\exists \sigma^{\text{CDi}} : \Psi^\mathcal{U} \rightarrow (\Psi^\mathcal{S})^*. (\sigma_\Psi(\psi, \gamma) \equiv \sigma^{\text{CDi}}(\psi))$  Thus,  $\sigma_\Psi$  maps  $\mathcal{U}$  commands to  $\mathcal{S}$  commands without considering the state in which they will be executed.

Since  $\mathcal{T} \stackrel{\text{sim}}{\underset{\text{CDM}_w}{\mathcal{S}}}$ , by weak model containment,  $\mathcal{S}$  commands are mapped to  $\mathcal{T}$  commands without considering the state in which they will be executed. Call this mapping  $\sigma^{\text{CDM}}$ .

Thus, let  $\sigma' : \Psi^\mathcal{U} \rightarrow (\Psi^\mathcal{T})^* = \sigma^{\text{CDM}} \circ \sigma^{\text{CDi}}$ , and say  $\sigma_\Psi(\psi^\mathcal{U}, \gamma^\mathcal{T}) \equiv \sigma'(\psi^\mathcal{U})$ . This forms a command mapping that satisfies CDi.

Hence,  $\mathcal{T} \stackrel{\text{sim}}{\underset{\mathcal{P}}{\mathcal{U}}}$ . ■

**Lemma 26** Given access control systems  $\mathcal{S}$  and  $\mathcal{T}$  and simulation properties  $\mathcal{P} = \{\text{SCa}, \text{QPa}, \text{CDi}, \text{R} \rightarrow\}$ ,  $\mathcal{S} \leq^\mathcal{P} \mathcal{T} \Rightarrow \mathcal{T} \stackrel{\text{sim}}{\underset{\text{CDM}_w}{\mathcal{S}}}$ . That is, if  $\mathcal{T}$  is at least as expressive as  $\mathcal{S}$  with respect to properties  $\mathcal{P}$ , then  $\mathcal{T}$  admits a CDM weak simulation of  $\mathcal{S}$ .

*Proof:* To prove this lemma, we let  $\mathcal{S}$  and  $\mathcal{T}$  be arbitrary access control systems such that  $\mathcal{S} \leq^\mathcal{P} \mathcal{T}$ , and we show that  $\mathcal{T} \stackrel{\text{sim}}{\underset{\text{CDM}_w}{\mathcal{S}}}$ .

Since  $\mathcal{S} \leq^\mathcal{P} \mathcal{T}$ , for any access control system  $\mathcal{U}$ , if  $\mathcal{S} \stackrel{\text{sim}}{\underset{\mathcal{P}}{\mathcal{U}}}$ , then  $\mathcal{T} \stackrel{\text{sim}}{\underset{\mathcal{P}}{\mathcal{U}}}$ .

Since  $\mathcal{S}$  can trivially simulate itself,  $\mathcal{S} \stackrel{\text{sim}}{\underset{\mathcal{P}}{\mathcal{S}}}$ , and thus  $\mathcal{T} \stackrel{\text{sim}}{\underset{\mathcal{P}}{\mathcal{S}}}$ .

Thus, given  $\gamma_0^\mathcal{S}, \gamma_1^\mathcal{S} \in \Gamma^\mathcal{S}, \gamma_0^\mathcal{T} \in \Gamma^\mathcal{T}$ , by SCa and  $\text{R} \rightarrow$ , if  $\gamma_0^\mathcal{S} \stackrel{a}{\sim} \gamma_0^\mathcal{T}$  and  $\gamma_0^\mathcal{S} \mapsto \gamma_1^\mathcal{S}$ , then

$$\exists \gamma_1^\mathcal{T}. (\gamma_0^\mathcal{T} \mapsto^* \gamma_1^\mathcal{T} \wedge \gamma_1^\mathcal{S} \stackrel{a}{\sim} \gamma_1^\mathcal{T})$$

By CDi,

$$\exists \sigma^{\text{CDi}} : \Psi^\mathcal{S} \rightarrow (\Psi^\mathcal{T})^*. (\sigma_\Psi(\psi, \gamma) \equiv \sigma^{\text{CDi}}(\psi))$$

This relation is thus a weak access-containment relation, which satisfies the definition for a CDM weak simulation, and hence  $\mathcal{T}$  admits a CDM weak simulation of  $\mathcal{S}$  ( $\mathcal{T} \stackrel{\text{sim}}{\underset{\text{CDM}_w}{\mathcal{S}}}$ ). ■

**Theorem 27**  $\text{CDM}_w \doteq \{\text{SCa}, \text{QPa}, \text{CDi}, \text{R} \rightarrow\}$ ; that is, the CDM weak simulation decomposes to authorization correspondence, authorization preservation, independent command mapping, and forward reachability.

*Proof:* By Lemma 25, if  $\mathcal{T} \stackrel{\text{sim}}{\underset{\text{CDM}_w}{\mathcal{S}}}$ , then  $\mathcal{S} \leq^\mathcal{P} \mathcal{T}$ . By Lemma 26, if  $\mathcal{S} \leq^\mathcal{P} \mathcal{T}$ , then  $\mathcal{T} \stackrel{\text{sim}}{\underset{\text{CDM}_w}{\mathcal{S}}}$ . Thus,  $\mathcal{S} \leq^\mathcal{P} \mathcal{T}$  if and only if  $\mathcal{T} \stackrel{\text{sim}}{\underset{\text{CDM}_w}{\mathcal{S}}}$ , and thus the CDM weak simulation decomposes to  $\{\text{SCa}, \text{QPa}, \text{CDi}, \text{R} \rightarrow\}$ . ■

### I. CDM Strong Simulation

For the purposes of this proof, let the set of simulation properties  $\mathcal{P} = \{\text{SCa}, \text{QPa}, \text{CDi}, \text{CS1}, \text{R} \rightarrow\}$ .

**Lemma 28** Given access control systems  $\mathcal{S}$ ,  $\mathcal{T}$ , and  $\mathcal{U}$ ,

$$\mathcal{T} \stackrel{\text{sim}}{\underset{\text{CDM}_s}{\mathcal{S}}} \wedge \mathcal{S} \stackrel{\text{sim}}{\underset{\mathcal{P}}{\mathcal{U}}} \Rightarrow \mathcal{T} \stackrel{\text{sim}}{\underset{\mathcal{P}}{\mathcal{U}}}$$

That is, if  $\mathcal{T}$  admits a CDM strong simulation of  $\mathcal{S}$ , and  $\mathcal{S}$  admits a simulation of  $\mathcal{U}$  with properties  $\{\text{SCa}, \text{QPa}, \text{CDi}, \text{CS1}, \text{R} \rightarrow\}$ , then  $\mathcal{T}$  admits a simulation of  $\mathcal{U}$  with properties  $\{\text{SCa}, \text{QPa}, \text{CDi}, \text{CS1}, \text{R} \rightarrow\}$ .

*Proof:* To prove this lemma, we let  $\mathcal{S}$ ,  $\mathcal{T}$ , and  $\mathcal{U}$  be access control systems such that  $\mathcal{T} \stackrel{\text{sim}}{\underset{\text{CDM}_s}{\mathcal{S}}}$  and  $\mathcal{S} \stackrel{\text{sim}}{\underset{\mathcal{P}}{\mathcal{U}}}$  but are otherwise arbitrary, and we show that  $\mathcal{T} \stackrel{\text{sim}}{\underset{\mathcal{P}}{\mathcal{U}}}$ .

Choose an arbitrary state  $\gamma_0^\mathcal{U} \in \Gamma^\mathcal{U}$  and command  $\psi^\mathcal{U} \in \Psi^\mathcal{U}$ , and let  $\text{next}(\gamma_0^\mathcal{U}, \psi^\mathcal{U}) = \gamma_1^\mathcal{U}$ . Let  $\gamma_0^\mathcal{S} \in \Gamma^\mathcal{S}$  such that  $\gamma_0^\mathcal{U} \stackrel{a}{\sim} \gamma_0^\mathcal{S}$ . Since  $\mathcal{S} \stackrel{\text{sim}}{\underset{\mathcal{P}}{\mathcal{U}}}$ ,

$$\exists \gamma_1^\mathcal{S} \in \Gamma^\mathcal{S}. (\text{terminal}(\gamma_0^\mathcal{S}, \sigma_\Psi(\psi^\mathcal{U}, \gamma_0^\mathcal{S})) = \gamma_1^\mathcal{S} \wedge \gamma_1^\mathcal{U} \stackrel{a}{\sim} \gamma_1^\mathcal{S})$$

Let  $\gamma_0^\mathcal{T} \in \Gamma^\mathcal{T}$  such that  $\gamma_0^\mathcal{S} \stackrel{a}{\sim} \gamma_0^\mathcal{T}$ . Since  $\mathcal{T} \stackrel{\text{sim}}{\underset{\text{CDM}_s}{\mathcal{S}}}$ ,

$$\exists \gamma_1^\mathcal{T} \in \Gamma^\mathcal{T}. (\gamma_0^\mathcal{T} \mapsto \gamma_1^\mathcal{T} \wedge \gamma_1^\mathcal{S} \stackrel{a}{\sim} \gamma_1^\mathcal{T})$$

Thus, there exists a  $\mathcal{T}$  command  $\psi_0^\mathcal{T}$  such that  $\text{next}(\gamma_0^\mathcal{T}, \psi_0^\mathcal{T}) = \gamma_1^\mathcal{T}$ . Define  $\sigma_\Psi : \Psi^\mathcal{U} \times \Gamma^\mathcal{T} \rightarrow \Psi^\mathcal{T}$  such that it returns  $\psi_0^\mathcal{T}$  for  $\gamma_0^\mathcal{T}, \psi^\mathcal{U}$ .

Then, given  $\gamma_0^\mathcal{U}, \gamma_1^\mathcal{U} \in \Gamma^\mathcal{U}, \gamma_0^\mathcal{T} \in \Gamma^\mathcal{T}, \psi^\mathcal{U} \in \Psi^\mathcal{U}$  such that  $\text{next}(\gamma_0^\mathcal{U}, \psi^\mathcal{U}) = \gamma_1^\mathcal{U}$ , and  $\gamma_0^\mathcal{U} \stackrel{a}{\sim} \gamma_0^\mathcal{T}$ ,

$$\exists \gamma_1^\mathcal{T} \in \Gamma^\mathcal{T}. (\text{next}(\gamma_0^\mathcal{T}, \sigma_\Psi(\psi, \gamma_0^\mathcal{T})) = \gamma_1^\mathcal{T} \wedge \gamma_1^\mathcal{S} \stackrel{a}{\sim} \gamma_1^\mathcal{T})$$

Hence,  $\mathcal{T} \stackrel{\text{sim}}{\underset{\{\text{SCa}, \text{R} \rightarrow\}}{\mathcal{U}}}$ . Next, we show QPa.

Choose an arbitrary request  $r \in \mathcal{R}^\mathcal{U}$  and states  $\gamma^\mathcal{S} \in \Gamma^\mathcal{S}, \gamma^\mathcal{T} \in \Gamma^\mathcal{T}$ . Since  $\mathcal{S} \stackrel{\text{sim}}{\underset{\mathcal{P}}{\mathcal{U}}}$ ,

$$\sigma_Q(r, \gamma^\mathcal{S}) \equiv \gamma^\mathcal{S} \vdash r$$

Since  $\mathcal{T} \stackrel{\text{sim}}{\underset{\text{CDM}_s}{\mathcal{S}}}$ ,

$$\forall r^\mathcal{S} \in \mathcal{R}^\mathcal{S} : \sigma_Q(r^\mathcal{S}, \gamma^\mathcal{T}) \equiv \gamma^\mathcal{T} \vdash r^\mathcal{S}$$

Since  $\mathcal{S} \underset{\mathcal{P}}{\text{sim}} \mathcal{U}$ , by SCa,  $r \in \mathcal{R}^{\mathcal{S}}$ . Thus,

$$\sigma_Q(r, \gamma^{\mathcal{T}}) \equiv \gamma^{\mathcal{T}} \vdash r$$

Hence,  $\mathcal{T} \underset{\{\text{SCa}, \text{QPa}, \text{R} \rightarrow\}}{\text{sim}} \mathcal{U}$ . Next we show CDi and CS1.

Since  $\mathcal{S} \underset{\mathcal{P}}{\text{sim}} \mathcal{U}$ , by CDi and CS1,  $\exists \sigma^{CDi} : \Psi^{\mathcal{U}} \rightarrow \Psi^{\mathcal{S}}. (\sigma_{\Psi}(\psi, \gamma) \equiv \sigma^{CDi}(\psi))$  Thus,  $\sigma_{\Psi}$  maps  $\mathcal{U}$  commands to  $\mathcal{S}$  commands without considering the state in which they will be executed.

Since  $\mathcal{T} \underset{CDMs}{\text{sim}} \mathcal{S}$ , by strong model containment,  $\mathcal{S}$  commands are mapped to single  $\mathcal{T}$  commands without considering the state in which they will be executed. Call this mapping  $\sigma^{CDM}$ .

Thus, let  $\sigma' : \Psi^{\mathcal{U}} \rightarrow \Psi^{\mathcal{T}} = \sigma^{CDM} \circ \sigma^{CDi}$ , and say  $\sigma_{\Psi}(\psi^{\mathcal{U}}, \gamma^{\mathcal{T}}) \equiv \sigma'(\psi^{\mathcal{U}})$ . This forms a command mapping that satisfies CDi and CS1.

Hence,  $\mathcal{T} \underset{\mathcal{P}}{\text{sim}} \mathcal{U}$ . ■

**Lemma 29** *Given access control systems  $\mathcal{S}$  and  $\mathcal{T}$  and simulation properties  $\mathcal{P} = \{\text{SCa}, \text{QPa}, \text{CDi}, \text{CS1}, \text{R} \rightarrow\}$ ,  $\mathcal{S} \leq^{\mathcal{P}} \mathcal{T} \Rightarrow \mathcal{T} \underset{CDMs}{\text{sim}} \mathcal{S}$ . That is, if  $\mathcal{T}$  is at least as expressive as  $\mathcal{S}$  with respect to properties  $\mathcal{P}$ , then  $\mathcal{T}$  admits a CDM strong simulation of  $\mathcal{S}$ .*

*Proof:* To prove this lemma, we let  $\mathcal{S}$  and  $\mathcal{T}$  be arbitrary access control systems such that  $\mathcal{S} \leq^{\mathcal{P}} \mathcal{T}$ , and we show that  $\mathcal{T} \underset{CDMs}{\text{sim}} \mathcal{S}$ .

Since  $\mathcal{S} \leq^{\mathcal{P}} \mathcal{T}$ , for any access control system  $\mathcal{U}$ , if  $\mathcal{S} \underset{\mathcal{P}}{\text{sim}} \mathcal{U}$ , then  $\mathcal{T} \underset{\mathcal{P}}{\text{sim}} \mathcal{U}$ .

Since  $\mathcal{S}$  can trivially simulate itself,  $\mathcal{S} \underset{\mathcal{P}}{\text{sim}} \mathcal{S}$ , and thus  $\mathcal{T} \underset{\mathcal{P}}{\text{sim}} \mathcal{S}$ .

Thus, given  $\gamma_0^{\mathcal{S}}, \gamma_1^{\mathcal{S}} \in \Gamma^{\mathcal{S}}, \gamma_0^{\mathcal{T}} \in \Gamma^{\mathcal{T}}$ , by SCa and  $\text{R} \rightarrow$ , if  $\gamma_0^{\mathcal{S}} \stackrel{a}{\sim} \gamma_0^{\mathcal{T}}$  and  $\gamma_0^{\mathcal{S}} \mapsto \gamma_1^{\mathcal{S}}$ , then

$$\exists \gamma_1^{\mathcal{T}}. (\gamma_0^{\mathcal{T}} \mapsto \gamma_1^{\mathcal{T}} \wedge \gamma_1^{\mathcal{S}} \stackrel{a}{\sim} \gamma_1^{\mathcal{T}})$$

By CDi and CS1,

$$\exists \sigma^{CDi} : \Psi^{\mathcal{S}} \rightarrow \Psi^{\mathcal{T}}. (\sigma_{\Psi}(\psi, \gamma) \equiv \sigma^{CDi}(\psi))$$

This relation is thus a strong access-containment relation, which satisfies the definition for a CDM strong simulation, and hence  $\mathcal{T}$  admits a CDM strong simulation of  $\mathcal{S}$  ( $\mathcal{T} \underset{CDMs}{\text{sim}} \mathcal{S}$ ). ■

**Theorem 30**  $CDMs \doteq \{\text{SCa}, \text{QPa}, \text{CDi}, \text{CS1}, \text{R} \rightarrow\}$ ; that is, the CDM strong simulation decomposes to authorization correspondence, authorization preservation, independent command mapping, lock-step, and forward reachability.

*Proof:* By Lemma 28, if  $\mathcal{T} \underset{CDMs}{\text{sim}} \mathcal{S}$ , then  $\mathcal{S} \leq^{\mathcal{P}} \mathcal{T}$ . By Lemma 29, if  $\mathcal{S} \leq^{\mathcal{P}} \mathcal{T}$ , then  $\mathcal{T} \underset{CDMs}{\text{sim}} \mathcal{S}$ . Thus,  $\mathcal{S} \leq^{\mathcal{P}} \mathcal{T}$  if and only if  $\mathcal{T} \underset{CDMs}{\text{sim}} \mathcal{S}$ , and thus the CDM strong simulation decomposes to  $\{\text{SCa}, \text{QPa}, \text{CDi}, \text{CS1}, \text{R} \rightarrow\}$ . ■